In the Matter of:

10th Annual FTC Microeconomics Conference

November 2, 2017
Day 1

Condensed Transcript with Word Index

For The Record, Inc.
(301) 870-8025 - www.ftrinc.net - (800) 921-5555

1

FEDERAL TRADE COMMISSION

THE TENTH ANNUAL

FEDERAL TRADE COMMISSION

MICROECONOMICS CONFERENCE

DAY 1

Thursday, November 2, 2017
9:00 a.m.

Federal Trade Commission
Washington, D.C.

2

FEDERAL TRADE COMMISSION

3

P R O C E E D I N G S

\- - - - -

WELCOME AND OPENING REMARKS

MR. VITA: My name is Mike Vita. I am the Deputy Director for Research and Management, as well as currently Acting Director here at the FTC's Bureau of Economics, and I just want to welcome everybody to the -- this is now the Tenth Annual Micro Conference we have held here, and it's hard to believe it's been that long.

The purpose of this conference, like all of our conferences, is an attempt to combine, you know, cutting-edge academic research with discussions of real-world policy problems, and I think if you look at the agenda, I think, you know, it promises to do that this year like it has in the past.

Before the first panel gets started, just a few announcements and a few acknowledgments. First of all, I want to express our gratitude to Northwestern University and the Searle Center for their continued cosponsorship of this conference.

Let me also acknowledge, you know, some of the work of the FTC people who -- first of all, the Scientific Committee that helped us put this together. That's Steve Berry from Yale, Jonathan Zinman from

4

Dartmouth, and Igal Hendel from Northwestern. Thanks. You know, as usual, really great work in selecting a great selection of papers to be presented.

Just a few words about us, I mean, and some of the people in the group -- in the audience might not be overly familiar with the FTC and what we do here. We're an independent agency that, along with the Department of Justice, enforces the antitrust laws, but we also have -- unlike Justice, we also have an additional enforcement mission, which is consumer protection law.

So these enforcement missions are supported by the FTC Bureau of Economics, which is about 80 Ph.D. economists, and so it makes it one of the largest applied microeconomics groups in the Federal Government. We think -- you know, those of us at the FTC think that these twin enforcement missions reinforce and complement each other. Competition is most effective when consumers are making well-informed decisions and free choices, and consumer protection works best when consumers have real alternatives.

So the purpose of today's conference, like its predecessors, is to help ensure that the FTC's actions are informed and guided by the best possible economic analysis. So in addition to the normal papers, you

5

know, the paper sessions -- which one will be starting in just a few minutes -- we also have two panel discussions that focus more on policy issues. The panel discussion today is antitrust-focused. It's on cross market hospital and healthcare provider mergers. Tomorrow, there will be a panel discussion on privacy and data security. So, you know, both of -- you know, those each addressing the two -- the twin enforcement missions of the FTC.

So I thanked our scientific panel, and let me also thank the FTC economists who, you know, helped organize this, Ted Rosenbaum and Nathan Wilson, and Peter Nguon of the Bureau of Economics, one of our RAs who really did great work in helping put this together. And also our admin team which works -- does incredibly hard work behind the scenes to make sure that this comes off. So Maria Villaflor, Kevin Richardson, Neal Reed, Constance Harrison, Priscilla Thompson, Tammy John, and Chrystal Meadows.

Before I turn this over to the -- to Jonathan Zinman and the first panel, let me call your attention to two calls for research that recently were issued by the FTC just in the last couple of days. The first is a request for empirical research and public comments on the effects of certificates of public advantage and

6

other kinds of state-based regulatory approaches intended to control healthcare prices and quality.

COPAs have turned out to be pretty important for the FTC, especially our hospital merger enforcement mission. Basically, if two hospitals that are -- you know, that are close rivals propose to merge and it would ordinarily attract the attention of possibly an enforcement action with the FTC, that can be avoided by obtaining something called a COPA from -- it's awarded by the individual state, and that can immunize the transaction from antitrust scrutiny, and that's come up in a couple of recent cases. So it's an important issue for us, and we would like to know more about you know, how these things work and what their effects are.

So if you go to the FTC's website and also, you know, out on the table where the papers are, you'll see the actual call for research. There's going to be a public workshop in 2018 where -- you know, where researchers can, you know, present the results of their findings, and we can, you know, help maybe, you know, guide further policy actions by the FTC. So, again, you can find a discussion of that -- you know, the call for papers is out, and there's a copy of it up on the table out there with the papers, but it's

7

also on the FTC's main website, on the homepage.

Second, Economic Inquiry has just announced a symposium on the economics of consumer protection. The goal of the symposium is to create a unique reference on consumer protection economics that would synthesize what we know about the current state of economic analysis, of consumer protection law, and enforcement policy, identify what consumer policy questions are in need of more analysis, and advance the application of economics to consumer protection policy analysis and law enforcement.

The symposium -- which there will actually be a symposium next year, next December here at the FTC -- celebrates the 40th anniversary of the 1978 founding of the Division of Consumer Protection in the Bureau of Economics. So up until 1978, the Bureau of Economics really only was directly involved in the antitrust enforcement mission. By the time the late seventies rolled around and the enforcement mission was picking up steam, it was realized, you know, that there needed to be a role for economics there, too. So next year is the 40th anniversary.

So I think there will be a special issue of the journal where selected papers are -- you know, are published and then the symposium in December. The

8

editors of the symposium are Wesley Wilson, he's one of the editors -- I guess he is the lead editor of Economic Inquiry -- and Jan Pappalardo, who's our Assistant Director for Consumer Protection here at the FTC. And, again, that call for papers, a copy of it's out at the desk, but I think it will be also posted on our website.

Okay, I think that's all the things I wanted to announce. Oh, just, you know, I am supposed to make announcements about exits and things like that. So if there is a fire or something, follow the exit sign. You guys are all Ph.D.s. I'm sure you can figure that out. There's a cafeteria over here -- there is going to be lunch, but there's a cafeteria over here if you want to get something to eat this morning, you know, you can go over there, and we also have coffee and other refreshments back there. And I think that is it.

So, Jonathan I think is running the first session. Is that -- is that right? Okay, you want to do the -- okay, so I will hand it over to Nathan Wilson. Thanks.

PAPER SESSION

MR. WILSON: Thanks a ton, Mike.

All right. Our first paper session will kick off with Charlie Murry speaking on middlemen as information intermediaries. Before Charlie takes over, I wanted to pass on one last important announcement, which is the restrooms are directly behind me. The men's room is on the left as you're facing this wall. The women's room is on the right. All right. And if you have questions, please just let us know.

Charlie?

MR. MURRY: Okay. So this paper has a title, and the title is on the first slide. So this paper is about middlemen or intermediaries. It's with Gary Biglaiser, who's here today, and Fei Li at UNC, and Yiyi Zhou who's at Stony Brook. So the next slide, please.

Okay. So middlemen, intermediaries are everywhere in the economy. Just kind of from a very broad perspective, there's kind of a public debate of whether middlemen or intermediaries are of value to society. So think about, like, you know, in different industries like financial institutions or different kind of used goods industries.

Our paper is going to be about used cars, okay, and we all might have a particular thought or vision about a used car intermediary or a used car dealer, all right? And so -- the next slide, yeah -- so this may or may not be your kind of picture of what a used car dealer is, but we're going to provide some framework to think about the services this type of intermediary provides in his marketplace. Okay, next slide.

Okay, so there's kind of two things that the literature is focused on for the role of intermediaries. The first is that intermediaries facilitate search and matching by potential buyers, okay? So there's a large literature, a theoretical literature on the role of intermediaries fulfilling search and matching. There is also a very nice empirical literature documenting this feature of different industries.

We're going to take a different view of the role of intermediaries. We're going to view intermediaries as being information certifiers, okay? So there is some theoretical work on the role of intermediaries as certifying information in markets -- Biglaiser '93 and Lizzeri '99 -- but there is very limited empirical work documenting or testing the role

of intermediaries of relieving informational problems and certifying goods in markets.

So what we're going to do is examine the role of used car dealers in relieving asymmetric information. So specifically we're going to present a model of dealer experts motivated by features of the used car market, and then we're going to empirically test two key assumptions of our model.

The first assumption is that -- it has to do with the role or the value of dealers in this market, okay? What we find from the model, what we predict from the model is that there's a price premium that the dealer can charge over the private party market, and so you can -- as a consumer, you can go to the dealer to buy a used car, right, or you can do an off-dealer transaction with an individual. And so there's a price premium that the dealers can charge in this market, and this price premium is correlated with the age of the car.

In particular, we find that the price premium is increasing in the age of the car, and the important thing about the age of the car is that it's a -- is the age of the car is correlated itself with the fact that the car might be a lemon or not, okay, and I am going to go over that in detail.

The second way we bring the model to the data is kind of more of a classical test of asymmetric information. So the model predicts that cars sold from private parties turn over more quickly, so they're resold more quickly than cars sold from dealers. That's because cars sold from private parties are more likely to be lemons, okay, and the people purchasing those cars are going to want to shed those cars, get rid of those cars.

Okay. So why do we care about this? So why do we care about used cars? Why do we care in general about this question? So there's kind of two reasons, big picture reasons. One is kind of from an academic perspective. Used cars are a very classic example of kind of Akerlof's lemons problem. In particular, we're suggesting that dealers here, as an information certifier, act as a counteracting institution in the parlance of Akerlof 1970, okay? So these guys are -- they come in and they make the market work.

More practically, why do we care about used cars or why do we care about this question? Well, there has been a lot of recent research on online markets and the role of information certification in online markets -- so, for example, you know, the star rating on eBay -- yet a significant amount of trade still happens offline. So how do these offline markets function without these kind of mechanisms that we've grown accustomed to in online markets, like a star rating, right? You cannot go to your friend who's selling a car and ask him for his star rating, right? He's never sold a car before. No one's rated him, okay?

In particular, the used car market is quite a relevant market when thinking about asymmetric information problems. So, first of all, it's a huge market. These numbers are kind of good guesses. We've done a lot of work to figure out how big the used car market is, but there is not a lot of great information. But the total sales of used cars in a year is roughly 300 to 400 billion dollars, okay? This is roughly three to four times the gross merchandise volume for eBay in a year. So this is a very large market.

Cars are kind of the classic example of asymmetric information. They're complicated machines that require specialized care. And so we think this market is ripe for asymmetries. In our sample, the sample of transactions we have, about two-thirds of used car transactions happen through a dealer, and there are kind of institutional features that we think

make this market kind of natural to study this kind of information certification problem.

That is that dealers are quite regulated by U.S. states. They might have reputation concerns that are different than private parties that are transacting in this market, and they do things to kind of explicitly try to resolve asymmetric information problems, which is offer warranties and guarantees.

Okay, so what I'm going to do is I'm going to present a model very briefly. I'm not going to use any notation, so I am just going to give you the intuition for our model and then give you the intuition for the predictions of the model. Then I'm going to bring the model to the data and show you how we test these two predictions with our data.

Okay, so let me set up the model here. So in the model we have different agents interacting. The first type is an owner of a car or a seller of a car, okay? The owner of a car has a used car. That car can have two potential states. So that car can either be high quality or low quality. The car can either be not a lemon or a lemon. This state of quality is private information to the owner of the car.

With some probability, a quality shock arrives -- this is a continuous time model, so a

quality shock to the car arrives at some rate that changes the state of the car, that takes the car from a high state to a low state, so a nonlemon to a lemon, okay? Also, there's a liquidity shock that arrives to the owner of the car so that the owner is forced to sell this car at some point in time.

When the owner receives this liquidity shock or the seller receives this liquidity shock, they can visit a dealer with some exogenous probability, okay? So they are basically allowed to visit a dealer with some probability, okay? So let me talk about what happens if they visit a dealer.

If the owner of the car visits a dealer, the dealer can run a test to discover the true quality of the car. So the dealer in this market is an expert, okay? Other private individuals in this market are not an expert, so they cannot run the same test. The dealer then makes a take-it-or-leave-it offer to the owner of the car, and this could potentially be a losing offer. For example, if the dealer finds out this is a lemon, then the dealer might not want to take possession of the car.

Okay, then if the dealer takes possession of the car, they set a selling price to the market, and they earn some profit, the selling price they set

minus the cost that they paid for the car from the original owner. And then also the dealer has some cost of selling a lemon, okay? So if the dealer takes possession of a lemon and decides to sell it on the market and sells it, there's some cost there, okay? And I'll talk more about that cost, but this makes it kind of -- this makes it so that the dealer has a distaste for selling a lemon to the market.

Okay, there's two more -- okay, there's one more agent in the model that's interacting with the dealers and the owners, and that is the buyer, okay? So we assume that there's at least two buyers for any given car in the market, and the buyer receives some -- like a single unit of utility from a car until it turns bad, okay? So they continue to receive this flow utility from the car until the car becomes a lemon, and then they receive no utility from the car.

The two or more than two buyers simultaneously bid on the car, whether -- if it's from a private market or whether it's from a dealer. And what does the buyer know? Okay, so the buyer observes the vintage of the car, observes how old the car is. This is important because we have this kind of -- this quality shock arriving at the car at some random rate, okay? So with an older car, it will more likely be a lemon, all right?

But the potential buyer does not observe the true quality of the car, and also, the potential buyer does not observe whether the original owner has taken the car to the dealer to be inspected and tried to sell the car to the dealer, okay? So the only thing this buyer knows is the age of the car.

Okay. And then there's one more thing that can happen in the model, that after this stage where these buyers bid on a car and potentially transactions happen, there's a second stage where the new owner of the car can resell the car, okay? So these new owners of the cars receive another liquidity shock with some probability, and when they receive that liquidity shock, they must sell the car.

Okay, if you're a new owner of the car, you can also just sell your car anyway. So, for example, if you took possession of a lemon and you realized it was a lemon, you can also go in the market and try to get rid of this lemon, okay? And in the resale market, it works very similar to the original market, and that is in the resale market, the new buyers observe the vintage of the car, how old the car is, but they do not know the selling motive of the new seller of the car. So they don't observe if the car was purchased

originally from a dealer or a private party, and they don't observe the private information of that car, what state it's in.

Okay, so there are some key assumptions we've made in this model, and I just want to briefly kind of go over them. One is that there's an exogenous liquidity shock, so basically I have a car and I have to sell it for some reason. Then I can only go to the dealer with a certain probability. What does this do in our model?

Well, this kind of forces there to be a mixing between high and low cars in the private market, okay? So the consumers will know for sure that there's some probability that a high car will exist in the private market. Okay, why is this important? Well, we need some sort of, in Akerlof's term, counteracting institution to kind of make this market go, okay? Otherwise, the consumers, all right, would believe that the only cars being sold are lemons, and the market would unravel like Akerlof 1970.

So another example of this is in a nice paper by Igal Hendel and Alessandro Lizzeri, where they basically tranche up the market into different -- into new cars and used cars, and they create distribution evaluations, and this is another kind of way to get

the market to go. So this is our mechanism to get this market to go.

We can actually allow -- the model is very simple if you don't allow for endogenous self-selection, but we can actually allow for some endogenous self-selection -- so for some guys who have a lemon to endogenously go to the market and sell this lemon -- but we need some high type cars in the market.

The second important assumption is the value of a high car versus the value of a low car. The key assumption here is that the utility that a consumer gets from a high car is greater than from a lemon. Both types' value depreciation with age, so as the car gets older, the flow utility you receive from a car depreciates, and the difference between a high car and a low car goes to zero as the age goes to infinity, okay? So at some point your car just is a POS, and it doesn't really matter if it's a lemon or not. You don't really want to be driving it.

Also, the third kind of assumption I want to point out is that the dealer has some sort of cost of selling a lemon. This is kind of a way to model the fact that dealers are less myopic than private sellers, okay? For example, dealers might have --

dealers are going to this market day-in and day-out, whereas private sellers go to this market typically once and don't go back, okay? And so dealers might be concerned about their reputation and so they might not want to sell a lemon for some reason.

Okay. In the model, buyers will bid their expected quality in the private market. A seller will accept a dealer's offer if it's greater than the outside option. It turns out that dealers will only trade in high types of cars, and the price that they set equals the buyer's utility, so the flow utility for this high type of car, and that the resale -- there is action in the resale market.

Okay. So the model predicts kind of two things that we're going to take to the data. One is that the dealer's price premium in price terms is humped shaped -- and I'll show you this in a second -- and that the dealer's price premium in percentage terms is greater than one. So there is always a dealer premium, and it's increasing in the age. So in percentage terms, there's an increasing dealer premium.

Okay. The intuition here is that older cars are more likely to be lemons. Buyers value the dealer's kind of certainty. Buyers value the fact

21

that the dealers are screening these cars and providing them certainty of whether this car is a lemon or not, okay? But as the car gets really old, the depreciation wins out, and really old cars are worth nothing anyway, and so that's why this is humped shaped in dollar terms.

The other prediction we're going to take to the data is that -- it has to do with car resales. So a buyer is less likely to resell a car if originally purchased from a dealer, okay? The intuition here is that all lemons, when -- if you are in this resale market and you've found a lemon, you want to get rid of it, okay? But the only way you are going to get rid of a car that's not a lemon in this resale market is if you receive a liquidity shock. So it's more likely for cars to be resold if they were bought originally in the private party market as opposed to the dealer market.

Okay. So to test these assumptions, we gathered data on the universe of used car transactions in the states of Pennsylvania and Virginia, and the nature of these data are the following: We have the transaction date, we have something about the vehicle identification number, we have the odometer reading, we have the zip code of the buyer, we have the seller

22

identity. So whether if it's a dealer or not, who the dealer is, and if it's a private party, we know the zip code of the individual.

For the Virginia data, which we used to test the dealer premium story, we observe a -- kind of a long panel, 2007 to 2014, and we observed a squish VIN, which is like the first 11 digits of the VIN. So we don't know the exact car being sold, but we know everything else about the car, so what -- the exact trim, the specifications of the car.

For the Pennsylvania data, we have a much shorter panel, so 2014 to 2016, but we observe the entire VIN of the car. So we know exactly which car is being sold, and so we can link the same car over time, subsequent resales.

Okay. So just to point out some moments from the data, this is from the Virginia data, it's clear that there is a dealer premium. Not conditioned on anything, there is a dealer premium. So, on average, the price of a private party transaction is about $4,000 in our data set. The price of a dealer transaction is about $13,000.

Dealer transactions are younger, six years as compared to 11 years, and they have lower mileage on them, okay? This is not surprising. About 60 percent

23

of our sales go through dealers, okay? So this is what the dealer premium looks like in the Virginia data. This top line is the average price -- is the price -- the average price of a dealer sale. This bottom line is the average price of a private party sale at different ages, okay? So it's clear that prices are going down as the age of a car gets greater, okay?

There is this hump -- you can't really see it here, but there is this humped shaped in the dollar terms of the dealer premium, so it starts out pretty small and then gets bigger, the difference between these two lines, and then goes down again. And then this line here is the dealer premium in terms of a ratio of the dealer price to -- the average dealer price to the average private party price.

Okay, but these patterns could exist because there's some kind of sorting in these markets based on observed characteristics of the car; for example, the make or model of the car, okay? So we are going to do a little bit more serious job about testing this idea that there's a dealer premium and that it's increasing in age.

So we are going to run a Hedonic price regression, and the important thing about this

24

regression is we're going to be able to add a make/model/MY/trim/fixed effect, so this is going to condition on basically everything observable in terms of characteristics about each car.

We are going to include a seller type and car age dummy interactions, okay, and so we're going to be able to predict the dealer premium for any given model/make/ MY/trim for any given age, okay? And I am going to show you what these -- the basically expectation of these prices look like, conditional on these controls, on the next page. I won't go into the different samples we used.

Here's the predictions from this kind of price Hedonics regression, okay? So this is the predictions of the -- this is the predicted price premium for a dealer by age, okay? And so you can see that young cars have a low price premium. For example, one-year-old cars have about a $1,000 price premium, on average, and it certainly is hump-shaped, and it peaks at about six years, okay, and then it depreciates, right? So this is consistent with the prediction of our model, that in dollar terms, the price premium for dealer cars is hump-shaped, okay, and depreciates as the car gets older.

Okay, the other prediction from the model about price premium is that, in ratio terms, it's increasing, and so this is the predictions in terms of the predicted price premium in terms of ratios from this kind of Hedonic pricing model, okay? And so for one-year-old cars, the price premium is about 15 percent, and this increases until about age eight or nine and then kind of levels off or slightly decreases to age 20. Okay, so these two features of the data are consistent with the predictions of the model about the dealer price premium.

The other thing we test is this bit about car resale. So the implication is that a buyer is less likely to resell his car if the car was purchased from a dealer. Okay, we used the Pennsylvania data where we can link cars over time, so we take all cars that we observed transacted in 2014 and follow them until 2016, and this is just the moments from the data here.

One percent of dealer sales are resold within a quarter, 2.2 of private sales are resold within a quarter, and these patterns continue to exist as we think about longer resale terms, so two quarters, three quarters, four quarters. So it's always the case that private sales are kind of more likely to turn over within any of these resale bins.

So, again, this could be because there's kind

of mixing and types of cars that are being resold, so we -- let's see, this guy -- can you move the slide deck one forward? Yeah, okay, so we do a little better job here. We look at resale rates by three-month intervals, controlling for the same model/make/MY/trim/fixed effect, and we're worried that there might be some reason -- unobserved to us -- why you buy a car from the dealer in the first place and you are going to sell it quickly. For example, maybe you're a transient person who's just in town for three months. So we are going to instrument for seller type by using data we have on the inventory holdings of dealers in local markets.

And we run a logit model with fixed effects, and when we do instrumenting, we use a control function to instrument for the -- whether the car was bought from a dealer or not, and our results here are consistent with the patterns in the data, which is that if the car was bought from a dealer, it's more likely to be resold in one quarter than a car bought from a -- I'm sorry, it's less likely to be resold after one quarter than a car bought from a private party, and in two quarters, and in three quarters, and in four quarters.

And actually, this kind of probability that you

resold within this time frame is decreasing in the time frame, which, you know, kind of is suggestive evidence that this asymmetric information problem is kind of going away over time. And this is the results with the instruments, but they tell the same story.

Okay, so my time is almost out. Kind of one important thing, though, that might be going on in this market that we haven't talked about, although I mentioned it briefly at the beginning, is that we could observe this dealer premium because of a search and matching role for dealers, and so in the paper we spend some time talking about the predictions of a search and matching story in our model, and we show that it's not quite consistent with these particular patterns that we find; specifically, the fact that the dealer premium is increasing in car age.

But I'm out of time, so I won't go over that in detail, and I will leave it at that. So, thank you again.

(Applause.)

MR. WILSON: Thanks very much.

Discussing this paper will be Tobias Salz of Columbia University.

MR. SALZ: Thank you.

So thanks to the organizers for allowing me to

discuss this paper. I like the topic. I like the paper. It was a lot of fun to think about it.

So I want to start out with this quote that I found, and I'm sure many people out here are familiar with the story behind this. The reason I thought it was fitting is that if you are sort of in the trenches of day-to-day work and, you know, you kind of look at other papers and you feel everything is pretty incremental and the process comes very slowly, and then you compare sort of how we nowadays think about the interplay between theory and data to -- you know, what this referee must have thought when he was rejecting the original lemons paper, you can't help but think that, well, actually, we have made a lot of progress, and I think this paper is a nice example of sort of how we use data in a more nuanced way to inform theory.

So as the authors have already highlighted, intermediation is a big part of the economy. I think there's more empirical work to be done. And something to appreciate about this paper is that it's really hard to pin down these informational stories without observing sort of what people know exactly, and what this paper does, it leverages the intertemporal dimension of this market a lot, and basically it gets a lot of qualitative prediction just out of that from a very, very simple model, okay?

The paper has great data. I think I learned many new facts. The model is parsimonious in a good way and yet it gives all these predictions. And then something else to highlight, I think I -- and I know we have studied car markets a lot, so it's one of those markets that, you know, we think we know very well, and yet there's a lot more to be learned here, which I think is nice, and it's a big and an important market.

So let me quickly recap the model. I'm using notation, so I should have coordinated a bit better, but -- so the model basically has cars that are aging, and everybody can condition on age, so there's a depreciation effect that everybody can condition on, and then there is a quality that's only -- a binary quality value that only the seller can condition on. And so over time the car becomes sort of obsolete just because it -- you know, it depreciates but also because the chance is getting higher that it becomes of low quality.

And the sellers of the car, they exogenously meet either a party in a bilateral market or a dealer, and dealers will only sell high-quality cars. So this

is -- in the middle, I think -- in the main text it's pretty much assumed, but then in the appendix, it's derived from a little cost that the dealer has to maintain reputation.

And so over time, the market -- the bilateral market becomes more and more select -- adversely selected because dealers, they are -- they are confronted with a car, they turn it down, and so more and more low-quality cars will be traded in the bilateral market.

So one thing that's also important which makes the model sort of much simpler is that the sellers in this market can basically extract all the rents, okay? So buyers engage in Bertrand competition for cars.

So I think that you can pretty much get three out of four predictions from the model by just looking at these two equations. So as I said, dealers will only sell high-quality cars, and they offer a warranty with these cars, so that buyers know that they get high-quality cars, and so dealers are charging the value -- the high utility value.

And in the bilateral market, you have this ratio of good cars, the mass of good cars that are being traded in the bilateral market, times the utility value for a high-quality car. The low-quality

value cancels out because it's normalized to zero.

And so what you can nicely see from this expression is that because over time the high-quality cars -- the mass of high-quality cars is shrinking, that the ratio of bilaterally traded cars is going to one. And then if you look at the ratio between the price and this -- what buyers paid in the bilateral market, you'll see that the depreciation effect, which comes through you, cancels out, and so you are left with the selection effect, and it's then pretty easy to see that the percentage premium in this market increases over time.

And then lastly, if you take the difference between these two, so if you subtract the bid for a car in the bilateral market from the price that a dealer is charging, you have both the depreciation effect and the selection effect, and this leads to this hump-shaped pattern that the authors are documenting.

And then lastly, there's an additional prediction which comes from an extension of the model in which sellers are able to resell.

So I want to make two main comments here. The first comment is that dealers and bilateral sellers are engaging in a slightly different business. So

dealers, they're negotiating with the customers over a bundle, over a bunch of products at the same time. They're negotiating over the car, financing and insurance, trade-in value, add-ons, and so on.

So one thing I was wondering is whether the car that we observe for dealer -- sorry, the price that we observe for dealer-traded cars is the price that you would get if you only negotiate over that part of the bundle, okay?

And so just to sort of give you one piece of evidence, this is borrowing something from a paper that I'm working on which is looking at how dealers price the financial aspect of the deal and the car price jointly. So what this shows here is a regression of prices of financial charges and the total price, which includes both financial charges and the car price, on a bunch of controls.

So, again, we have model controls and a bunch of other controls for the buyer, and then the key variable of interest here is an indicator that's called subvented. So this is whether or not a loan is subvented. So what is a subvented loan? It's basically when the vertically integrated lender of the manufacturer -- so Honda Finance -- ties the dealer's hands and says, well, you have to offer this loan as zero percent finance, okay?

And so this is -- you know, you should take this with a grain of salt, because this is not -- you know, subvented is not randomly assigned here, but I think there's some descriptive evidence from these subvented loans that this is a joint pricing problem and that, in fact, as you can see, the car price goes up conditional on the model if the dealer can no longer get this markup from the financial aspect of the deal, okay?

And so, of course, mechanically, financial charges go down, and it actually turns out that the total price is going up for these subvented deals. So this is just saying, well, there might be some other aspect of the bundle that is not included in this price here, and so I'm not even sure I would necessarily go against the authors, because it depends a lot on sort of, you know, where in the age distribution of cars are financial services offered and sort of how are these relative markups assigned, but I think one thing to sort of dig into a bit more is sort of what -- you know, what kind of price are we looking at here?

So then the other thing that I wanted to explore a bit more -- and this is sort of going back

to I think a debate that, you know, is a long debate in the schooling literature. You know, it's either all information and selection and signaling, or it could be, you know, added value, and here it could be that dealers are, in fact, you know, adding value or recovering some of the value of low-quality cars.

And so I'm proposing a very simple model. This is exactly -- pretty much exactly like the one that the authors are looking at except that now the quality of the car is also observable to everyone, and dealers can recoup some of the lost value of a low-quality car at some fixed cost.

So they pay a mechanic for a few hours, and then there's sort of a random shock with which they can recover some of the value, and then I'm sort of playing around with the fixed cost for this -- for the mechanic, and on top of that, you might think there is a fixed cost for inventory, for dealer inventory, okay? Then we can kind of see how, over the time of or over the age distribution, these patterns that the authors are documenting look like, okay?

So I am going through a few cases now. So the first case is where there's a repair cost but no fixed cost. So the dealers basically repair the car if the fixed cost of the repair is smaller than the value

that they can recover. And what you can see is we basically are totally wrong on the market share. The market share of dealers is flat because they're always selling all cars, because quality is observed, right, and they can extract all the rents.

Then we get that both the percentage premium and the absolute premium has this hump-shaped pattern. So basically we get one out of the three prediction of the model without the extension, right, so that's not going to give it to us. Ah, here we go.

So now the case where we have a fixed cost but not a repair cost. So what you can see here is that now dealers are turning away cars that no longer -- whose value is no longer larger than their fixed cost of holding, so over time they are actually losing market share because they are more likely to send these low-quality cars back to the bilateral market.

We also see that the dealer percentage premium is increasing, but the absolute dealer premium is also increasing, and this comes due to selection on varieties. So at the very end of the age distribution, you only want to hold expensive varieties. And so, again, we get this time two out of three right, and it's not going to give it to us.

Okay. So now both of these patterns combined

give you sort of a similar picture. We have decreasing market share, we have increase in percentage premium, but we also have increasing dollar premium. So the only case that I could find -- the only case that gets all three of those patterns right in the unextended model is the following case, and you can sort of think for yourself whether you think that's a plausible model.

So dealers have this repair cost, and they take a low-quality car only if they can repair it, and they send it back to the bilateral market otherwise, okay? Remember, this is a bit of a funny -- this would be a bit of a funny case because quality is observable. So I could still sell the car at the low value, so there must be some other reason, and so maybe I don't want to have shabby-looking cars on my lot.

So this might be one reason I send -- send away cars that are of -- that I can't repair, and then you get basically all three of those patterns, and I would just basically urge the authors to sort of maybe discuss a bit more what can be done with this alternative. I think a plausible explanation is that the dealer is actually adding some value.

So then I have a few more other smaller comments. I think the model in this resale extension

37

makes also a sharp prediction on whom resales go to; namely, they should not go to dealers, because sellers know that dealers know the quality. So there should basically be an additional prediction that the model makes, that the authors can test for.

This is maybe bickering a bit, but, like, you could also look at more types than -- at least in these equations that I showed you, these -- the -- you don't -- you no longer get the depreciation effect, so cancel out nicely. I think things are still going through, but that would be something to look for.

Then I was wondering about spatial controls, so where are dealers located and does this correlate with the types of cars that are being sold in some way that could give rise to some of these patterns.

In terms of model specification, you could think that sometimes the bilateral market gives you all these guarantees or warranties that the dealers in this market are providing. So, for example, if I'm selling to my brother-in-law, then, you know, I don't want to sell him a lemon necessarily, so there might be some sort of repeated play that enforces this or has this reputation effect.

So something I was wondering about, what about age versus mileage? So basically the model is all

38

done in terms of the age of the car, but you could think that -- you could sort of rephrase all of this and say, well, the observable dimension is actually the mileage of the car. We instead control for age, and I was wondering whether all of these things sort of look similar if we instead do it the other way around.

And then something I'm always interested in is sort of the distribution of prices. So we see these are all mean effects, but, you know, if you sort of see the price distribution for dealers in the bilateral market, how do they look like? Is it driven by the tails? That would be something to look at.

And that's all I have.

MR. WILSON: Thanks very much. We have got time for a couple of questions for Charlie. If you do have a question, please wait for the microphone to assist our stenographer.

Charlie, do you want to come up?

Jonathan?

MR. ZINMAN: This may be a question for FTC folks as much as for Charlie. I'm wondering if there are any efforts under way, public sector and/or private sector, to bring more data to bear in this market against the asymmetric information problem. It

39

seems like there's a lot of data that could, in principle, be captured about vehicle history.

So, for example, it seems like many cars these days know what's wrong with them on an automated basis. Is there any interest in helping or forcing manufacturers to capture and share that data to build vehicle histories that would be more observable?

MR. MURRY: So I can't really answer that question too well. One resource available to kind of researchers -- not in the Government -- there's more research sources, like, for example, CARFAX reports. Those are, in practice, accessible to researchers. The problem is CARFAX reports are not very good actually, so you shouldn't really trust a CARFAX report.

And so maybe there is a role to kind of regulate CARFAX reports or something else, but I can't -- yeah, I don't know if anybody from the FTC wants to take that or not.

MALE AUDIENCE MEMBER: Thank you for that paper. I really wish my good friend and colleague Jim Lacko were here, because I don't know if you're aware of his research -- it's 30 years old now -- but he had research from a survey of used car buyers, and I would just encourage you to look back at his paper from, I

40

think, 1985, where one of the differences with his data from your data was he was able to get more information on the private sales -- was it to a brother-in-law, as the discussant mentioned, or was it to a stranger -- to see whether that factor influenced the quality of the car being traded.

I don't know his research well enough to do more than that, but I would just encourage you to look back at that report. We'll be happy to get it to you. I'm here at the FTC, but it's a really, really fine piece of work with a different -- a different approach to the data.

MR. MURRY: Yeah, so one thing to kind of -- on that point and something that Tobias mentioned, we do see the zip codes of the buyer and the seller in the private party transactions, and so we do have one specification that maybe not was in the paper that you saw, where we control for the zip codes of the buyer and seller.

And you might think in rural areas there might be a reputation effect of selling to somebody in the same zip code, but in urban areas, there isn't. So we -- but thank you for the suggestion, yeah.

FEMALE AUDIENCE MEMBER: I wanted to follow up on Tim's comment and question because I was thinking along the same lines. In Jim's work, I believe that he found that there was a distinction between the outcome if the dealer sold new cars and used cars or only used cars, and I'm wondering if you were able to investigate that.

MR. MURRY: Yeah, we do know the -- who the dealer is, so we do -- we do the same analysis for just the subset of dealers who also sell new cars and just -- and for the subset of dealers who only sell used cars. And the patterns that I showed exist for both types of sellers, but they're kind of shifted up for the used-only cars, but we do some of that in the paper, where we've split the sample into these two types of sellers, yeah. Thank you.

MR. WILSON: Thanks very much.

Now we're on to our next paper by Maryam Saaedi of Carnegie Mellon. She will be speaking about certification, reputation, and entry.

MS. SAEEDI: Okay. So this is a paper with my student, Xiang Hui, who is on the market now, Giancarlo Spagnolo, and Steve Tadelis. So in other sessions that we just saw, there exist in many markets, online and offline. If you want to buy something on eBay, you -- the seller will know more about the item that you know. If you want to get

something on Airbnb, you know, the host knows much more about the noise level in the neighborhood, and actually, as I found out last night, this is true for the hotels as well. There might be a train outside your hotel that goes every 15 minutes from 5:00 a.m. So I have been awake since 5:00.

And then, like, if you want to hire someone on Upwork, they know much better about their knowledge, their experience, or even if there are offline markets, you are hiring a procurement contractor, they know much more about the quality of their work than you do. And we know that from Akerlof, that there can be a lot of inefficiency and a lot of low-quality trade or sellers in the market as a result of that.

So there is a common solution for this problem, is having a reputation mechanism. So eBay, since its site has these feedback rating and other system has started, there are like Better Business Bureau, there are like restaurant ratings, Yelp reviews, all different kind of feedback ratings that can help overcome this problem.

So here in this paper we are actually focusing more on other kind of mechanisms, not exactly, but something that is related to feedback rating that can mitigate some of this asymmetric information problem.

Okay.

So one standard solution that is similar to what we just saw is certification. So the certification can be that marketplace can be using data or some kind of process to certify the quality of their existing sellers. So the problem with doing this kind of licensing is that it can be some kind of barrier for the new sellers to enter in the market.

And this kind of certification is actually very common in online markets. So eBay has eBay Top-Rated Seller. Airbnb has super -- Airbnb Superhost. Upwork has its top-rated freelancers.

So these badges sort of show that there is -- shows to the buyers that these sellers have passed some bar. And, for example, on eBay, when you're searching for something -- so this is from the time that we have the data on -- we're searching at that time when iPod still was traded heavily on eBay, and when you were searching on eBay, you would see that there is, like, some sellers that are top-rated, so -- and you could see that on the search page, and then after you click on them, you would see more information on the listing page about the fact that this seller is top-rated and more information about this seller as well.

Okay. So the good thing about the badges is that it mitigates some of the asymmetric information, but the problem is that it can be a barrier to entry. And what we want to do here in this paper is, what will happen if you actually make this certification to be harder to get? You want to see what's the incentive for the new sellers entering into the market and what is the quality distribution of these sellers and sellers in the market.

And we are going to use a study -- a policy change on eBay to answer these questions. So I will skip the literature review. So I will give you some stylized model that can help you think about what we have in mind. It's a very simple model. It's actually based on a paper that I have with Ugal Oppenheim (phonetic)here.

So we are assuming here that -- the paper is -- so sellers are competing in a competitive market. So it's quite -- not a very bad assumption for eBay. There are hundreds of sellers, if not thousands, that are selling the same product. So we are assuming that it's a competitive market.

And then firms differ in two dimensions, either in their quality or entry cost. So their quality, they're assuming they can have three levels of quality here, z1, z2, and z3, and they have an entry cost that is independently distributed from a function G.

And the buyers observe certification badge. They care about the quality, but they don't see z1, z2, or z3. They see only the certification badge, and the certification badge signals if the quality is weakly above a threshold.

So we are assuming that there is a baseline demand function, which is P(Q), which would be the demand for the lowest quality, and the demand for equality with expected quality, z-bar, is going to be additive to that demand function. So it would be P(Q) plus z-bar.

So the policy change is going to be having this format. So at the beginning, before the policy change, group of z2 and z3 sellers were getting badge, they -- and they would show that they have a badge, but afterwards, only z3 sellers would show that they have a badge. So we change the threshold from z2 to z3.

Okay, so the impact on the entry depends on the changes in the prices. So we can prove that the price for z2 guys is going to drop, so these are the people who were getting the badge before, now they are not getting the badge. And they are losing the premium,

and as a result, the price will go down, and there would be fewer of these sellers in the market afterwards.

And -- but for z3 and z1, we can't prove that the price would go up for them for sure, but we can prove that at least one of the prices go up, but it can be both of them. So if, for example, for z3 types, if the price is going up, it's because now they can get a more informative signal that they are of the highest level of quality, and then they would be entering into the market more.

And for z1 type, if the price is going up for them, it's because now they are pooled with z2 guys, and they can get the higher prices, and they would be entering into the market more.

Okay. So I am now going to the data. So we have proprietary data from eBay, and we have a lot of information about all the transactions that happen and what has happened afterwards. So we see everything that the buyer can see, and we can see everything about the history of the seller and the history of the buyers.

And one thing about the eBay product catalog that we are using is that eBay has this catalog formation that is about 400-plus categories that will

include everything that is sold on eBay. So one category is fiction and literature, another one is fresh-cut flowers, and I will explain how we are going to use that.

Also, we have these product IDs that we are going to be using that is looking at very homogenous goods, like iPhone 6, black, 32-gigabyte, unlocked. So it would be very specific about the product that is sold. And we will have information about when the sellers enter the market in each of these categories.

Okay. So what was the policy change? So eBay used to have another badge called eBay Powerseller, and they have changed that badge and made it harder to get. So nowadays, if you want to become a badge, which is now called eBay Top-Rated Seller, you have to meet all the requirements for Powerseller, and then you have to meet some additional requirements.

And here -- and also, then, you cannot see if someone has a Powerseller status but not eBay Top-Rated Seller. So the only thing that you can see is the new badge. You don't -- the previous badge is completely obsolete.

Okay. So the impact on the percentage of people, percentage of sellers who were badged after this change was quite stark. So about 10 percent of

the sellers had badge before, but afterwards, it dropped to about 4 percent, and then you see there is a growing in the number of sellers who have badge afterwards.

Okay. So what is our empirical strategy? Okay, so we can't be just using what has happened and just state the averages and say that's the impact of this policy, given that there's many things that are going on on eBay and also the fact that it's -- we are in the middle of financial crisis when this change has happened.

So what we are going to be using or doing -- we are going to do a two-stage approach. The first stage is we are going to be looking at the categories that I mentioned, the 400-plus categories that we have, and we are going to see which of them were more impacted by the policy than the others.

So here we run this simple regression, and we are looking at the share of badged sellers in each category over time, and we had a dummy for policy and some fixed effects and some time trend. And so this identification is based on assuming that these different markets were affected differentially, and it was exogenous why they were affected differentially.

We would run a placebo test to make sure that this doesn't just show some correlation that is driving other results as well. I will go over that by the end of the presentation.

Okay. So the second stage, we are going to use the results from the first stage and then look at different variables of interest, like number of entrants, quality and performance of entrants, and also quality of incumbents, and see if they were affected more with the policy or not. So we are multiplying this Gamma by that Beta-hat-C, and then we do also some other controls.

Okay. So the first stage result, we can see that this is the Beta-C, see how different categories were affected. So this is showing the whole 400 of them, just writing a few of the numbers. You can see that almost, other than one category, everything else had fewer badge sellers, few badge sellers afterwards, and you can see that the effect is very different from one category to the other. We have a good distribution, variation in the effects of the policy in these categories.

Okay. And now let's look at the results for the second stage. So now we are using that Beta-hat-C, so this Beta-hat-C is negative, so more affected categories have a bigger negative number. So

Gamma less than zero means that there would be more entrants in more affected categories.

So here the Y, the first Y we have is entrant ratio, so it's the number of entrants at time t divided by the number of sellers at time t minus 1. And what we can see here is that there are more entrants into the more affected categories, and -- if we are looking at three months before and after, six months before and after, but if we are looking at the months seven to twelve afterwards, we don't see a significant impact. It seems that the entry happens very early on for the first three to six months, and then it doesn't, at least, continue as much as we move on.

Okay. So then we want to also see what -- the impact on the quality of the entrants. To look at the quality of the entrants, instead of looking at feedback ratings, which is usually 100 percent positive for all the sellers, we are using this effective positive rating measure, which is based on the paper by Chris Nosko and Steve, that they are looking at the number of positive feedbacks over the number of transactions instead of the number of feedbacks received, and they show in their paper that that's much better measure of quality, and they can

see what buyers do afterwards, and they can say this sort of predicts how happy buyers are about the transaction, than just looking at the feedback positive rating.

And what we see here is that there is a positive impact, so there is -- the more affected categories have, on average, better quality entrants into the market. If you are looking at six-month window, twelve-month window, and also -- which is plus/minus three, plus/minus six, and actually, for this one, even if we are looking for longer time periods, we still see an impact on more higher quality entrants entering into those markets.

So this study shows us the average -- that on average higher quality entrants entering into this market, but you might also ask, so, what about the distribution of entrants? So that's what we want to do now. So we want to see how was the distribution of the entrants.

So to do that, we are going to divide the entrants in each of these subcategories into deciles based on their EPP in the first year after their entry. And then for each decile, we will run this regression again.

So for each decile, we will consider -- so this

would be their Betas for that categories, but then for each decile, we are going to look at what was the impact of the EPP of that decile. This is the result for that Gammas. So this decile one is the lowest quality item. Decile ten is the highest quality sellers.

So here, the decile ten, which is the highest quality entrants, we have a negative coefficient, which means that they're higher EPP in more affected markets. So it sort of shows you that the -- if you were considering the distribution of the entrants, if you are looking at the highest decile, it has moved more to the right -- yes, okay, that's right, to the right.

And on the decile one, which is the lowest quality entrants, a positive coefficient -- we have a positive coefficient. Even though it's not significant, it shows that the letter EPP on the more affected markets, so the other tail, the left tail, also move more to the left.

So we have higher raise in the tails and a bit less in the middle. So it seems that even though we get average higher quality, it's sort of coming from the very highest decile, which makes sense, because those are the only people who have a chance of getting a badge. Not everyone can get a badge. It's a very small percentage of people who can get a badge.

Okay. So we also -- so this was strange effects for the incumbents -- for entrants. We wanted to see what's the impact on incumbents, because it also can tell you that maybe these are not better sellers who are entering into these categories. It's just that these sellers who enter, after they enter, they start acting better because of the change that has happened in the market.

So we want to see what's the impact on the sellers who stay -- who were there before -- before the policy, and we actually -- when we are looking at these EPP measures, we don't see much of impact right after the policy change. So this is -- zero is when the policy change has happened. So the blue one, the blue circles are showing the average EPP for the entrants who entered in this month, this month, and so on.

You can see that the average EPP for the entrants afterwards is much higher than before, but when you are looking at incumbents, there's not that much of, at least, noticeable change. And we did different kind of slices of the data. So here we are looking at incumbents in top EPP quartiles and various

different kind of cut, and we don't see much of impact.

So here we are looking at -- the blue one is the year of the policy, the green is the year after, and the red is the year before, and we don't see much of impact for -- if we are looking at different quartiles. And we also run regressions, we don't see much of impact for incumbents.

So the best that we could come up why that has happened is that maybe the sellers on eBay are doing the best that they can already and there's not that much room for improvements for them, but we don't -- we were surprised that we didn't find any results here. So that was at least surprising for us.

Okay. So we also look at the impact on prices. So we are looking at group -- so BB is the sellers that are badged before and after; BN, sellers are badged before and not after; not badged before and badged after. So this NB group is surprising for many people. So the thing is that on eBay, they check the badge requirements once in a month, so it might have been that you were getting your badge no matter what, and even now that the badge becomes harder, you're still getting a badge. So that's why you have some people who don't have a badge but have badge after.

And then you have people who didn't have badge before and after.

So we are looking at relative prices, so -- because you can have a lot of stories about sellers after they become badged, they are going to sell better items. So we are going to look at the listing price divided by product value, and then we find the product value to be the average price of the product in the posted price format for that product ID. So the product ID, for example, again, would be iPhone 6, 64 gigabyte, black, unlocked. It would be very specific.

And we also look at sales probability, so what's the chance that they can sell an item and what's the number of items they can sell and what was their market share. So, in general, this is what we find if we combine everything together, that the best -- so the guys that were not badged and after -- even though there are not that many of them, they are the most affected, obviously, but then you have the people who were badged before and after, they have a positive impact, and then sellers who were not badged before and after, they also see a positive impact, and these people who lost their badge, they see a negative impact.

And so these are the regressions for different -- so I will skip that, it will take a long time, and for the last four minutes I talk a little bit about this placebo test that we run.

So this is a big concern that maybe our result is driven by some serially correlated subcategory heterogeneity that is simultaneously correlated with the Beta-hat-C and Y, the variable of interest and number of entrants, the quality of entrants, and so on.

So we were -- so if we assume that this kind of correlation would persist over time, so if it says something about these categories that will have more entrants or higher quality entrants coming to their market, we should be able to predict the number of entrants if we are going to look at the number of entrants or quality of entrants two months, three months, a year before, and so on.

So we don't -- we have done this for different time periods, for three months, six months, and a year before, to just looking at September, that was the policy change year, and we run all these regressions one more time for the number of entrants, the quality of entrants, and so on, and none of the variables that we find is statistically significant.

And it's not proved that this is -- that this is not a problem, but at least it's reassuring that if you are looking at some other time period, you don't see any correlations, only at the policy time that you see some correlation happening in other categories. Okay.

So another interesting thing here to show that is also related to our model is that -- so the entrants that we looked at, we sort of lumped two different type of entrants into one. So if a seller for the first time starts selling in a new category that they haven't been selling at before, we consider them to be a new seller, but then you can also have very brand new sellers who were not on eBay at all and then they entered into the market.

So in these regressions, we are going to divide them. So we are calling this new sellers versus existing sellers who were entering into the market and also enter into eBay completely or entering into new categories. So the result that we find is actually very interesting. So here, when we are looking at the number of entrants, we see that -- so both of the -- so the signs are all the same, so they are all negative, and that's what we had for -- when we were combining them together, but the magnitudes are

different, and actually, they -- it speaks to our model.

So when we are looking at these numbers, these numbers are smaller than this, so it sort of shows that for the new sellers, the very brand new sellers, it's not as easy to enter even when they changed the policy change, but when you are looking at the effects on their quality, you can see that these numbers are much bigger than these. So you can think of that fixed cost of entry into eBay is higher if you are a new -- a completely brand new seller, but entering into new categories, it still has a fixed cost, but it's not as high.

So you would see more entry for these guys who have lower fixed costs, but on average, even though they are higher quality, they are not as high quality as these guys, which they would be entering with smaller numbers, but when they enter, they have much higher qualities, on average. Okay.

So I -- we run a bunch of other robustness checks, look at the percentile of Beta, different kind of -- looking into the -- looking at shares, looking at the numbers, and everything -- the signs of everything will stay the same, so -- I don't have much more time -- and also for the prices, we're looking at

differing kind of -- and also, we're looking at exit behavior. What we see is that this BN group, people shrink in the size, and they exit the market with higher percentage.

And that's -- that's -- ah, okay, here. So this is what we have done. So that was our question, how does more demanding certification affect entry? We find that we will get more entrants into the market and higher quality with fatter tails, and quality change for -- from -- mostly from improved selection. Not much has changed in the behavior of sellers.

And it has some kind of implication for digital platforms, so the -- this certification can impact the rate and quality of the entrants, and -- but the other finding that we have is that they can impact quality mostly through selection and behavior of the sellers.

Thank you.

(Applause.)

MR. WILSON: Thanks very much.

Our discussant is going to be Peter Newberry of Penn State.

MR. NEWBERRY: All right, thanks.

So it's good to be here. Twelve years ago I was an RA at the FTC, and I helped organize I think what was the prequel of this conference, which was a

conference on Ecommerce with Chris Adams. So it's good to be here. I enjoyed reading the paper. It was fun.

So, yeah, as far as big picture goes, both these papers that we've seen are motivated by information problems. So we know Akerlof tells us about the lemons problem, where if sellers or if buyers can't -- or if consumers can't identify low- and high-quality sellers, then only the lowest quality goods will sell.

So we had an example of used cars, but the kind of rise of e-commerce has made us think a little bit even more about this, because presumably the information asymmetries are worse in online markets because you can't touch the goods and you can't really interact with the sellers as easily as you can in offline markets.

So we have seen some institutions that are introduced to try to help with these problems. We have warranties and seller guarantees or return policies, dynamic reputation, and what both these papers have focused on is certification, right? So -- but what Maryam and coauthors do is think about this also as a barrier to entry, and specifically, if you think about dynamic reputation and certification in these markets, this is going to be a barrier to entry as -- especially if you think about dynamic reputation, it's hard to get stars and recommendations without selling anything, right?

So -- and when you first arrive to the market, everyone might think you're bad and you're not going to sell anything, right? Okay. So we're going to think about these institutions as -- or at least certification as also a barrier to entry.

Okay. So what does this paper do? I would say the way I think about it, what are the long-run effects of introducing or changing a certification program, and specifically on eBay? So if we think about entry, so the trade-off here kind of is do the incentives from higher prices for sellers outweigh these barriers to entry? So we will see more entry. And then what happens to the distribution of quality? How does overall quality change?

And when you think about the entrants, like, are there higher or lower quality entrants? And what happens to the incumbents? And then they also think about prices and market shares.

So the strategy, as we saw, is to utilize a policy change that occurred on eBay in 2009 that actually made certification more difficult for the

sellers to reach, and evidence suggests that this policy had heterogenous impact across product categories. So we see that stricter certification, qualifications led to more entry. This entry was from the top and the bottom of the quality distribution, and we saw this result that it didn't seem like incumbents were responding.

Okay. So my opinions, what do I like? So this paper is really well motivated. I think reputation mechanisms are really important in these markets, especially as they continue to grow. So solving these information problems are very important. There's clear policy implications here for these platforms. How should we organize the platform as to solve these problems, right?

And I like the idea -- a lot of papers -- empirical papers on information asymmetries are looking at are information asymmetries a problem, you know, how do consumers react to this information, where here we're thinking about more of an equilibrium entry model, which I really appreciate.

And the data is obviously great. I mean, you have proprietary data from eBay. You observe everything, basically. Then you have this nice -- this nice policy change that occurred on eBay.

All right. Where I think there's some more work to be done, so the model I know is very stylized and it's from another paper, but I would have preferred it be more posed as a puzzle, when does entry increase and when -- like, under what conditions would we see it actually decrease when those middle guys, right -- and, you know, under what conditions will we see quality change and quality not change? So I would have preferred -- maybe that's in the other paper.

I also -- this assumption of the perfect competition, I -- you know, I buy it, but at the same time, there's still price dispersion for the same products on eBay, I'm guessing, so where -- you know, where does that come from? And you never really talk about exit in the paper, so I'm wondering -- you know, in the model, you could think about exit, and even in the data.

The results, can we say something about what happens to concentration? Like, you look at market shares of individual sellers but not really how the market overall is concentrated before and after the policy. You talk about prices but never really, like, how -- you know, overall price levels, what's going to affect -- how is this going to affect consumers? Is

it making consumers better off? But it could also make them worse off.

And then eBay's incentives, like, you know, what are their -- are they trying to align incentives between them and the sellers or -- you know, what's the impact of eBay's bottom line on these platforms?

So what I am going to focus on for the rest is this empirical strategy. So I'm still worried about -- I know about the placebo test, but I'm still worried about the exogeneity of your instrument, and I'll talk about that.

Okay. So the primary analysis uses this -- this is the second stage, right, where you have the outcome on left-hand side and then the policy on the right, interacted with some measure of what I'm going to call exposure to the policy, right? So the intuition here is it's kind of a continuous treatment, where if you're more exposed to the policy, that's the treatment group, and the less exposed groups are the control.

So this is, I think, you know, related to what's called a Bartik instrument, which is, you know, you basically are -- have this interaction with a policy variable on how exposed -- it's in the labor -- it started in the labor literature, but a really good example is in this QJE paper that's looking at the effect of Cash for Clunkers on sales for cars, and they actually use -- in a local market, the exposure is how many clunkers are there in that market when the policy -- the day the policy gets enacted, okay? So this is kind of an ex ante measure of exposure to the policy, okay?

So a key assumption -- obviously we know this. So where do I think the policy -- where do I think the problem is? Okay, so in order to calculate their exposure, they run -- you know, they run this regression, where the Beta-hat is their measure of exposure, but my -- I think what's going on here is the problem is this is actually an ex post measure of exposure. After the policy happened, how did -- you know, which categories were more affected, right? So this is an ex post measure of exposure rather than an ex ante measure of exposure.

So, for example, so share badged is actually an equilibrium outcome which is a function of your left-hand side variable. So, for example, just if you think the change in the share badged simply could just be written as this, so the change in the share badged is a function of entry, right? If this category saw more entrants, that's actually going to change the

share badged. So the result of this is actually this is a mechanical relationship between entry and the policy -- the policy -- the policy estimate that you're estimating.

Okay. However, so, fortunately, I think this is actually solved pretty easily. Think about this Cash for Clunkers paper. My suggestion is to use a measure of ex post exposure in a given category. So on the day the policy was enacted, how many sellers would have become badged that day, right? So this is a measure of exposure within a given category.

You could also just determine categories ex ante yourself and say this category is probably affected more because it has more high-volume sellers or the quality of the goods may be less salient, so they sell more new and used goods, and so you could, ex ante, just choose categories that you think are better control groups and then use maybe that first regression as evidence that that's true.

And then you could also -- I know you said this maybe isn't great, but you could just take an event study approach and assume that the policy was exogenous and then see, within a category, you know, what happened to -- what happened to entry and quality.

Okay. So I have one minute left. So that's my main comment, and I think, you know, try these things and see what happens. I just have a couple other things.

So other suggestions, could you think about the effect of the policy on other signals of quality? So are sellers reacting in some other way, like are they showing more photographs? Are they -- is their description a lot bigger, the guys who don't get badged?

Are they changing their products within a category? Are they selling more new versus used goods? Like I said earlier, what happened to overall price levels? And then concentration.

Another one of your main results is this quality dispersion, and I worry that that's also somewhat mechanical, maybe not completely, but my suggestion here is why not just run your definitive estimation on some measure of quality dispersion, like the distribution of -- like the variants of quality or some, you know, measure of distribution of quality, rather than break these guys up into these bins, right?

Yeah, so I just have, like, random other thoughts that were supposed to be or that we'll talk

about offline and I'll send you, but those are my main -- my main -- my main comments. So thanks for listening.

MR. WILSON: Thanks very much. And, again, we have a few minutes for questions from the audience.

MALE AUDIENCE MEMBER: Hi. I realize there's a public face and maybe a private face, but what is eBay's public justification for increasing the sort of stringency of the certification?

MS. SAEEDI: So I guess -- so I wasn't in any kind of committees who are deciding these things, and it's very hard -- we tried to actually get them to answer for this. We were not successful in finding what they actually thought about. So something that they told us, like, they found is that the number of people who were badged were too many. They just wanted to reduce that. And a lot of times, they never -- actually, they never go back to see what was the result of what they have done. They usually see that -- so they give some benefits to people who are badged, and they were thinking that the money that they have to spend on that is too much, and they wanted to reduce the number of people who have the benefit. Yeah.

MALE AUDIENCE MEMBER: I don't know if you have already done a little bit of this, so this is either a question or a suggestion, but I think the -- what Peter suggested along the lines of looking at different categories and characteristics of different categories is just super interesting, both for the identification purposes he's suggested, but also just kind of validating the theory testing --

MS. SAEEDI: Right.

MALE AUDIENCE MEMBER: -- and also thinking about, you know, future policy implications, if they were to do a -- you know, a kind of more refined sort of policy that was more targeted.

MS. SAEEDI: Yes. So I will be talking with you guys afterwards, so say exactly what was -- so one thing along the -- the lines of what Peter was saying, so what we have done, we've looked at, for example, very short period of one week before and after to find out -- so we don't have that much entrants or exiters during that time, but your suggestion is great.

We will -- we can just look at the sellers who were active in months before and see how many of them would have lost their badge or not, and we can just look at all their, like, qualification, the way that eBay decides for them if they are going to get badge -- be badged or not, use that, and that would be

much cleaner instrument. Yeah.

MALE AUDIENCE MEMBER: Yeah. So fascinating study. I'm wondering -- and this comes back a little bit to Andrew's question -- to what extent any of you guys know what was happening in the market generally, because you have incredibly rich data that comes from inside a single firm, but, of course, there are lots of places one could buy an unlocked 32-gig black iPhone --

MS. SAEEDI: Right.

MALE AUDIENCE MEMBER: -- kind of thing. And in thinking about policy, in particular, but even management, right, it would be really helpful at least to have some context for what's happening in the market as a whole. We just don't know what the equilibrium is.

MS. SAEEDI: Yeah. So when we are talking -- like, I guess, what you are -- you mean is the dynamics across platforms, and unfortunately, we don't have data on what is happening outside eBay, but that's very important. Actually, a lot of policies that eBay is applying is a result -- like, in response to what Amazon is doing or other type of platforms are doing. But, yeah, we have to think about that, see if we can find some kind of connection, yeah. That's a

good...

MS. JIN: And for a platform like eBay, the count of sellers is very different from the volume or profit contribution from each seller --

MS. SAEEDI: Right.

MS. JIN: -- and we know the power sellers contribute a lot more to the platform. So have you tried to look at other angles, like sort of the quantity they sell or the fees eBay can get from those sellers? And that -- I would imagine that probably will give a different picture.

MS. SAEEDI: So when we -- so I went through that slide for one second. So we looked at the market share of the sellers, and what we see is that the sellers who have stayed badged increased their market share, but the sellers who lost their badge, who were at the top before, and they're not -- they lost their market share.

So that's another question, I think, Peter also suggested, so what's the impact on the total quantity that is sold on eBay, so that can be -- given that one group has become bigger, one group becomes smaller, that can have different implications. We have a -- we don't have that in the paper now, but, yeah, that's a good point.

So we -- in the theory papers that I have with Ugal, we are looking at what's the optimal threshold to be put to increase -- to maximize the total quantity sold, but we don't look at the (inaudible) population.

MR. WILSON: Thanks, everybody. We are now moving on to our third paper in this session. That's going to be by Matthew Mitchell of the University of Toronto.

MR. MITCHELL: Okay. It's great to be here to talk about a theory paper that I think is directly relevant to some important FTC policy. So this is sort of a paper about Twitter, so it is sort of a theory of people who make recommendations on Twitter. One of those two people is a meaningful recommender on Twitter. The other one is not, really.

So, broadly, you know, I'm interested in intermediaries, because there's a lot of stuff to consume out there, and it's pretty hard to know what to consume, so you are going to need to ask somebody their opinion. And so there are a lot of ways to go on the Internet and get some opinions about what you ought to consume, okay?

Now, I'm going to be focused mostly on a narrow topic, which is people that get advice on the internet frequently from something like Twitter, but where the person giving the advice is getting, at least sometimes, compensated by someone who wants some advertisement. So the reason that the title of the paper is "Free (Ad)vice," with "Ad" in parentheses there, is because, you know, the advice has some ads snuck in there.

That's actually a relevant policy issue at the FTC right now. People at the FTC sort of know that. How people on the internet should be required to disclose their financial relationships with the things they recommend is actually an issue that's led to -- I used to give this talk and I used to say, well, the FTC has given some guidance, but it's not obvious that they've actually gone out and fined people yet, but now they have. So the FTC is actually taking action on this policy, and I guess I sort of think of this as broadly related to ideas like fake news.

So let me give you an example of what I have in mind here. This is a recommender or an influencer on Twitter named Kim France. Kim France is a -- was a very successful fashion journalist. I'm using her as an example largely because she had a very successful career, which she essentially quit to do this instead because you can make more money doing this. And so

one of the things she does is sort of really explicitly links people to her own blog where she then gets a commission for things that people buy through her blog.

I use her also as an example because she has a nice "Frequently Asked Questions" where she just explains what it is that she does, and she explains that, you know, I do sometimes get a commission from things I sell on my blog, but, of course, she says I'd never, ever link to anything that I wouldn't think was a good idea for you to consume, okay? So she's bundling these links with her recommendations of what would be good things to consume.

Her Twitter feed contains lots of stuff. Here's another example of something she said on Twitter. "TheUnbreakableKimmySchmit is a miracle." Now, I don't know if that's an ad or not. It actually turns out -- I did a little digging -- and I think The Unbreakable Kimmy Schmit was doing a viral ad campaign on Twitter during that period. I'm not trying to get Kim France into trouble with the FTC, that's a pretty old tweet, but it's not obvious if that is advice or an ad for The Unbreakable Kimmy Schmit, okay?

So one of the things we want to ask is, if that is an ad, should she have to disclose that she's

taking money from the TV show? In case you're wondering, that's a TV show on Netflix. There's so many TV shows now that you might not even know all of the TV shows, which is why you need to go to Kim France and figure that out.

Now, so that's -- most of the paper, I'm going to talk about this, I think, key pillar of the FTC's mandate, as we were talking about. I do think it's related to the FTC's other mandate, which is competition policy. Here's Google. They give advice on the internet, and this is like the classic picture of someone Googling Trip Advisor recommendations but getting Google's recommendations instead, okay? So that -- the fact that Google has some market power there might be relevant, but I am going to talk more about the Twitter-type examples today.

So I'm going to just think about a simple model capable of understanding the basic trade-offs, and in the middle, there's going to be sort of a question. Why do you pay attention to these people on the internet? The answer is going to be because they have an incentive to build up a good reputation by giving you some pieces of good advice. That is, if everything Kim France ever said was useless and she was just taking money, you'd stop following Kim

France, okay? So it's going to be sort of a pure reputational model.

The model is going to have a lot in common because there are no transactions of money in the model between you, the follower, and the influencer, because that's usually the way these things work. You don't pay Kim France directly for what Kim France has to say. And as a result, the model is going to look a lot like these models from the recent contracting literature without monetary transfers, especially -- there's a lot of such papers, and I am not going to talk a lot about the literature, but I do want to specifically point out Li, et al., and DeMarzo and Fishman, which are sort of the two most closely related models to what I'm going to talk about today.

And so the thing I want to stress that's going on here about ads and that's going to be relevant to what I'm going to have to say about disclosure policy is that ads are sort of playing two roles here if you're a consumer. On the one hand, in the current instant, ads may be a temptation for the influencer to bias their advice away from what's best for you and towards what makes them money, but if you weren't worried about that, then there would be nothing to worry about with influencers taking money on the internet.

On the other hand, the fact that these influencers can make money by running ads is the way you encourage them to give good advice now. The fact that they will be able to run ads at future instants is the way they want to keep you around, and keeping you around is the motivation for giving good advice at the current instant.

So ads are really serving two functions here. They're not just a temptation for the influencer. They're also the way the follower gets incentives from the influencer to behave by sticking around for future ads. Because of that, it's going to turn out the disclosure is not unambiguously good here, and I'm going to propose that in this model there's an idea that's unambiguously better, which I am going to describe in more detail, which is going to be something I am going to call opt-in disclosure, where an influencer can decide and has to state whether or not they're living by, in some sense, the FTC's disclosure rules or not.

Okay. So I'm just going to tell you about the basic model. The idea here is to keep the model as stark at possible. So this is going to be a continuous time model with an infinite horizon,

blah-blah-blah. I am going to try to keep the notation as limited as possible, so I am going to normalize the discount rate to one, so these are forward-looking people with some discounting that I'm normalizing.

There's going to be two people here, a follower and an influencer. The follower decides whether or not to follow the influencer. That variable is going to be called f. That's going to be observable. And then good advice arrives to someone who's following an influencer at a rate Lambda times one minus the ad level a. So the idea here is the more is the underlying advertisement level of the advice, the less likely is it to generate good advice for the follower.

Now, in the basic model, there's a direct trade-off. If a is set to its maximum, which I'm taking to be one, there's no good advice. I have an extension where good advice can show up even when you're running the ad technology at maximum, but I want to keep things as stark here as possible to understand how this model works.

So the influencer is going to privately choose the ad level. The follower merely observes when they receive good advice. When the follower receives -- gets good advice, they get a value -- like a lump of

one unit of value. The influencer gets value Lambda times a for an ad level a. So in any given instant, the total surplus from this relationship is Lambda, okay? A just decides how that is split up.

I'm going to also show you explicitly -- this one I'll get to -- I'm going to explicitly show you what happens if ads generate waste in some sense, that ads lower the total amount of surplus, but for my benchmark model, they don't, okay?

The follower has an outside option if they decide not to follow, that's s. That's what makes following costly. If you decide to follow this person and they're not giving you any good advice, then that's leading to some cost for the follower that you could think of as s. Lambda is bigger than s, so it's better to follow than not if you're getting good advice.

I just want to point out, since Lambda is bigger than s, if we had full information here, we could just trace out the full information Pareto frontier. That would just be all the combinations of the follower's value and the influencer's value. I'm just getting out a little notation here. V is the follower's value, W is the influencer's value, where V plus W equals Lambda, okay? But, of course, that's

not what I want to study. I want to study the frontier under the asymmetric information I laid out in the last slide, where the level of the ad is unobserved to the follower.

So I'm going to describe this like a dynamic contract. That's not totally critical here, but it's, I think, going to make the construction as simple as possible. So there is going to be no monetary transfers here. The reward comes by a history-dependent choice of f and a. Think of that history as a complicated object. It's all the previous periods when you received good advice and whether you followed or not.

Of course, I am going to assume that the influencer can't commit to the actions they are going to take in this contract. I am going to assume for the purposes of this talk that the follower can commit to such a sequence of actions. That qualitatively doesn't affect the results here at all. It just makes the math a little simpler. Then I'm going to assume the influencer needs a fixed level to be willing to engage in being an influencer in the first place. This is sort of like a supply of influencers, okay?

So that contract's a potentially complicated object. It turns out it can be summarized by something simple, which is at any instant a forward-looking variable that describes in the contract for how long the follower will follow the influencer in expected discounted terms. I'm going to call that variable d. It lives between zero and one, because I normalized the interest rate to one, following forever, would mean a d of one. Giving up and never following would be a d of zero. Durations in between reflect different degrees, in some sense, of satisfaction with the influencer, because you are going to pay attention to them for longer on path.

It turns out that describing contracts this way -- and not as a function of the total history -- is without loss of generality, but if you want to think of this as just a restriction on the contracting space for now, that's fine.

I want to get to the -- to how this model works and then describe a little bit about policy. So the reason why this variable is very helpful is that the total surplus generated by this relationship is a simple linear function of duration. When these two parties are together, they get Lambda divided some way depending on the choice of a. And when they're apart, well, then, the follower gets their outside option, s, and the influencer gets nothing. So that's what that

says up there. The payoff to the influencer, W, plus the follower, V, adds up to, when they're together, Lambda, and when they're apart, s.

And so the fundamental idea here is that the bigger d grows, the harder it is to get incentives for the influencer to keep the ad level low, because the reason the influencer keeps the ad level low is to try to have these arrivals of good advice so that you're convinced to stay in the contract as a follower.

I mean, if you want to think about that at the extreme case, if d is equal to zero, surplus is as low as it can be, because we're going to have s forever, but the follower gets all of that. If d is equal to one, there will never be any good advice ever again, total surplus is as high as possible, but it all goes to the influencer. So the follower faces a tension between how much surplus they get and the total surplus in the relationship.

The way I'm going to think about this, like I said, is just like a contract. So let's think about incentive compatibility of a certain level of a. So the benefit from choosing a is -- I wish I could -- there is no way to point here, is there? It says that the marginal return to a is Lambda minus some stuff. The Lambda is the direct benefit of running the ad. I

assume you make revenue equal to Lambda times a from running the ad.

Whenever good advice arrives, duration, this d variable is going to change from what it currently is, d, to some future level -- I am going to call it d', and the value for the influencer changes from W(d) to W(d'). The more ads you run, the less likely that is to happen at rate Lambda. So that Lambda times W(d') minus W(d) is saying, when you run more ads, it's less likely that good advice arrives and the follower is happier with you.

So if the follower wants to get any good advice, they have to pay, in terms of future value, at least one unit. The amount that the influencer gets after giving good advice has to be one unit higher if they are going to not run ads. That's the incentive constraint here.

I'm going to show you what the value function looks like as a function of d, and then I'm going to explain to you why, and then I'm going to talk briefly about policy, and then I'm going to be done.

Here's the value function as a function of d. I have drawn as a function of d the dotted line. That's the total surplus in this relationship. The value function, of course, has to be below that. It's

a concave function. I'm not going to describe to you a proof of that here.

Of course, I said it starts out at the value being equal to s. It's a concave function where for a while it's strictly concave. That's the region where the advertiser is not running ads, the influencer is not running ads, and then it has a region at the top where the influencer is running ads.

So in this model, if you want to think of it, to the right we have influencers that have been more successful. They have given out good advice, and, therefore, this duration variable has jumped up to the right until we're in the region all the way to the right where they become a top influencer and stop running ads.

During that period, the duration variable is going to start to run downwards because they're not giving as much good advice -- in my benchmark model, they are not giving any good advice -- and for a while the duration variable runs down. They live off their reputation for a while, and after they live off their reputation for a while, we move back into the regime where a equals zero and good advice starts flowing again.

Again, that's a really extreme version of the benchmark model I want to show you, but the key feature that's true in a lot of the versions of the model that I do in the paper is these cycles for influencers. Influencers build up a reputation -- think of the low d as a sort of mediocre reputation -- they build their reputation by not running so many ads, and then they reap in the future from that reputation by running ads when d gets sufficiently high, okay?

I want to show you just a tiny bit more about how this contract works. Because V is concave, the incentive constraint it turns out has to bind, and the incentive constraint binding, you know in problems like this, is sort of fundamental to getting things well understood.

Let's think about what the incentive constraint binding means. It means, from the previous slide, that the amount by which W as to go up when there's good advice is exactly one unit. The amount that the follower receives every time there's good advice is exactly one unit. So it's as if, in future value, you're paying the influencer for exactly the value of the piece of good advice you received today, and you, as the follower, receive all the change in total surplus from the contract moving from d to d'. Total

surplus is increasing, and d' is bigger than d, so that's a positive number.

So the only reason we're ever in this range where a equals one is because you have not enough duration left, not enough d left to offer the influencer to possibly convince them to give you any good advice. Their reputation is so good that they have nothing to lose by running ads in the top range. In other words, we could characterize exactly that kink point, where you go from a equals zero to q equals one, that's exactly where the influencer's value is exactly one unit less than the maximum it could possibly be, and the maximum it could possibly be is Lambda.

Okay. Now I want to do some sort of policy-related experiments with that model. So in the model I assumed that a doesn't affect total surplus, but let's suppose it does. For instance, suppose that the return to the ad technology, instead of being Lambda times a, was Tau times Lambda times a.

Nothing about this math assumes that Tau is a number less than one, but I think that's the idea you want to have in your head, which is that perhaps running ads generates some inefficiency, some loss here, okay? And I want to characterize the contract

where, instead of Tau being equal to one in what I showed you up until now, Tau is a different number, like think of it as less than one.

As a function of d, this changes nothing about the contract. In particular, the allocation in terms of the choices of f and a is independent of Tau. The only thing that changes is that the influencer gets Tau-less, because when they run ads, the payoff is lower. What's the intuition here? Well, when you make ads have a lower return, you're doing two things: you're lowering the current temptation to run ads, and you're lowering the payoff in the future from any ads you might run if, as an influencer, you build up a good enough reputation to start running ads. You're doing those two things in exactly the same proportion.

So what I want to stress about this is pure taxes on ads here -- because that's another way you can interpret Tau, is a tax on ads -- have no effect on the amount of good advice that occurs in this model. If I was going to give a -- like, you know, I only get 25 minutes, so if I wanted, like, one piece of intuition from this model that's kind of different from a static model, it's that, because the dynamic effect of the taxes is exactly offsetting the static effect of the taxes.

So now I want to think about the FTC in this model. Igal has seen this paper before because I gave it at his birthday conference, and, of course, there were a lot of Minnesota guys there, and the Minnesota guys said, FTC? I don't know what that is. They wanted to know, what is the FTC in the model? You know, I have already got optimal contracts here. What do I need with the FTC?

So the way I'm thinking about the FTC in the model is that they have an additional technology that's not available to these two parties, a sort of auditing technology where they can go look, and if ads are run that are not disclosed -- I am going to describe a little bit sort of what I mean by "disclosed" -- then they can potentially punish someone who has chosen an a without disclosing that they're choosing, you know, that level of a.

So it's important that -- I'm assuming, of course, that the FTC has the access to some technology that these parties don't. Otherwise, use of that technology would already be incorporated -- already be incorporated in my optimal contract.

Okay. So here's how I'm going to think about disclosure rules by the FTC. I'm just going to think of them like a comparative static on the ad return in my model. So, first, suppose that the -- that the contract is calling for a equals zero. Any ads that are run there are a deviation from what the contract proposes.

Of course, let me just -- that deviation is relevant because that's what the incentive constraint is guarding against. So what I'm going to do is I'm going to make the FTC -- I'm just going to go back to my benchmark model where the ad technology, you know, starts from a return of one. They can make the return from those deviation ads lower, call it u less than one, okay?

Now, I'm going to assume that when a equals one, now the ads are on path, and I'm going to give the influencer a choice between whether or not they want to disclose or not disclose those ads. If they don't disclose those ads, the FTC is coming for them, so the payoff from those ads is u. If they do disclose those ads, I am going to allow for the possibility that those disclosed ads have a lower return, m, partially because influencers always say they do.

Also, because there's papers in the economics literature, like Inderst and Ottaviani, that say these kind of disclosure rules can lower the total pie

available, the total amount of surplus available to the two parties, the advisor and the advisee, like my influencer and follower. So I am going to allow for the possibility that the disclosure rules have costs like that, and I am going to think about what disclosure rules that potentially have costs like that do in this model. And then I'm going to show you what a disclosure rule would look like that would work better in this model, okay?

So, first, if the disclosure rules are weak, so that u is a number less than one but not as low as m, then nobody discloses any ads because, after all, they'd rather get m -- sorry, get u than get m by disclosing the ads. In that case, we know from our taxation results that disclosure is just a pure taxation on the influencers. It has no benefit for the followers.

On the other hand, if the disclosure rules are strict, meaning u is a smaller number than m, then they strictly benefit the followers because they make the incentive constraint easier to be satisfied. It shifts out the value function like that. Of course, what that means is as a function of u, welfare is not monotone, and I can't even tell you whether it's higher at the left end or the -- or the right end.

In the paper, I go into more detail to talk about where this lower return, m, might come from from disclosure, and the place I go is exactly the idea that there might be some ads that are also good advice. Kim France might sometimes get a commission from selling you a bracelet that she also thinks looks good and that you will think looks good, too, and I can write down a more specified model of disclosure where those disclosures can lead to costs because followers pay less attention to those particular recommendations. But, of course, in 25 minutes, I don't have time to do all that.

I just want to do one more thing in my last 50 seconds, which is describe what the model says is a better disclosure rule. The better disclosure rule here is what I'm going to call opt-in disclosure. So think of this as an influencer can decide -- just say on their Twitter bio, they could just say, "I follow all the FTC disclosure rules," or not say that.

Top influencers would want to opt out because they're in the reap portion of their cycle, and people who want to build a reputation would want to opt in because no one would pay any attention to them if they didn't. So in the model, that kind of opt-in policy is better than just a pure -- what I might call a pure

disclosure policy, and the reason is because of -- and I don't have time to talk about extensions -- and the reason is because sort of the fundamental difference here is that all the reward for influencers is coming from future payoffs. And so a way to tighten up the temptation to run ads when you're building your reputation, while still leaving the reward as high as possible when you've built a good reputation, is generally an improvement.

(Applause.)

MR. WILSON: Thanks very much. Our discussant will be Ginger Jin.

MS. JIN: Well, thanks for the opportunity to come back. My time at FTC must give staff impression that I can read theoretical papers, so they send me a real one to test out. The challenge is very much appreciated. I do wish that I had taken the graduate theoretical course more seriously 20 years ago, but I'm very grateful that Professor Mitchell has been very patient and responsive to my emails.

So this is a very interesting paper. What I like most is that it provides a novel framework that applies to both antitrust and consumer protection. Those in FTC know that we -- FTC actually run antitrust and consumer protection separately with very little overlap, but this theory sort of -- it's creative for us to think about search engine as an influencer providing advice to search engine users, just like social media influencer try to provide advice to Twitter followers. So I think that's a creative framework.

Actually, this framework could apply to any advertising-backed media, right? The radio, the magazines, television, if not all their income, most all of their income actually coming from advertising, and they think about the content they provide in order to generate followers. So I think in that sense this framework is very general.

Academically, it also naturally extend a lot of the literature in reputation, in paid advice, in disclosure, in the theory of market power. I would add to this list the theory of two-sided markets, as well as media bias.

Okay. So I just want to highlight the main insights in the basic model and probably give a comment on a few policy implications here. The basic model has five assumptions. The first is that influencer engage in an activity that's sort of disliked by the follower; namely, this advertising, okay? So here we assume away the influencer can

generate nonadvertising content that could be useful to the follower. I think that's a useful simplification, but extending along that direction might be interesting.

The second assumption is that the follower can only use following as the tool to generate reputation for the influencer. So the follower cannot say I'm going to pay more to a good advice if you have a good history or something like that. So it's unlike the typical reputation return, that you can get a return from higher price, and here sort of you can only get a return from the following behavior, and that following behavior is based on a noisy signal, which is the random arrival of good advice.

And following is costly, as Professor Mitchell said. It's because there's an outside option, so you can think of that as a potential competition with this technology here. The technology itself, that technology is exogenously given, okay? And that sets the total surplus to be fixed. So the tension in the basic model is how the influencer and the follower divide the pie rather than how to create a bigger pie.

So with those assumptions, the trade-off in front of the influencer is basically the trade-off between today and tomorrow. So today there is a pie

for you to grab, which is the advertising revenue, okay, and that pie might be small if you just have fewer followers, but it could be really big if you have a lot of followers, okay? So you can grab this pie today and leave nothing to the followers, or you can sort of keep the pie on the table and that's going to generate future good advice to the follower, and then the follower can decide what to follow or not, which determines tomorrow's pie, okay? So you're trading off between getting more of today's pie or leave it on the table and generating a bigger tomorrow's pie.

What's interesting here is that today's follower is actually going to affect the size of both today's pie and tomorrow's pie, okay? If you have a lot of followers today, today's pie is very big, but given that you already have a lot of followers, having a little advertising going on does not necessarily completely drain your follower crowd immediately tomorrow, okay? So your tomorrow's pie still depends on the good history you have generated so far, plus some not so good history in a day after today. So that's the trade-off in front of the influencer.

And as a result, we sort of see a cycling behavior. Okay, so I would call it sow and harvest.

That's the same as what Professor Mitchell call sow and reap. So from the influencer's point of view, over time, this reputation indeed is going up and down. In the down period, the influencer would have incentive just to, okay, I am going to refrain from advertising and just the sow the seeds and give good advice, and once I build up the reputation, I will be in the harvest mode, okay? I am going to harvest the advertising income because I have a lot of followers. That means today's pie is pretty big, okay?

From the follower's point of view, the follower sort of foresee the cycling behavior, but the follower can only have following as the tool. So the follower would say, okay, I am going to tolerate it with the harvest, because that's going to generate incentive for you to provide the sowing of good advice, but I am only going to tolerate it to some extent. If it's so bad, I am going to quit. I am going to quit forever, okay?

And that permanent quit is going to be a threat to the influencer, and with that threat, the influencer will not have incentive to overharvest, okay? So in the equilibrium, you are going to see this up and down, but the follower would follow. But the threat of equilibrium past will be important to

97

ensure that there is a sowing period before the harvest, okay? So in this sense, the harvest is sort of providing the incentive for the sowing, so the harvest is not necessarily bad thing, okay?

Okay. So with that insight, let me talk about policy implications. Before we get into the exact policy implications, I want to clarify what's the objective function we're looking at here. So are we looking at the follower's payoff as the objective function or are we looking at the total payoff, okay? I think that the position the paper takes is we put more weight on the follower's payoff.

In FTC language, that's -- we're maximizing consumer welfare rather than we're maximizing total welfare. The basic model set the total pie fixed, so it's just a redistribution question. The extended model probably can sort of vary the size of that pie. So I am going to focus in my discussion assuming that we are going to maximize the follower's payoff, okay?

Okay. So there are several tools to do that. So you could change the ad technology, including sort of the size of pie as well as the rule that's dividing the pie, or you can also restrict the influencer's behavior directly, like you cannot advertise or you have to advertise under certain rules, such as

98

disclosure, or you can raise the follower's outside option, which is kind of competing with this influencer in their good advice decision.

Okay. So one main result from the paper, arguing that advertising tax is neutral, the logic is that advertising tax is going to affect today's pie and tomorrow's pie proportionately, and your trade-off is between the two in a relative term, so it shouldn't matter. The tax should not matter because it's proportional; however, there is a fixed outside option there which does not go up or down with this tax. So my intuition is that when you have a lot of really high outside options, you would require a lot of good advice and expectation in this market before you follow, and that should generate incentive for the influencer to sort of restrain himself from harvesting to a greater extent and provide more good advice.

So my intuition is that this may not be completely neutral, because the -- the outside option is fixed, and then you change the advertising return, which would change the relative trade-off between that. So my hunch is that it may not be neutral in some contexts. So it will be good to see, and maybe I'm wrong.

Another extension is so far the model does not

99

allow the influencer to create nonadvertising content, right? But in a lot of social media examples, we see that they actually create entertaining videos or some opinion in Twitter, and that's -- that requires some effort to do, okay? So it will be interesting extension to see that what if there is a cost to create those nonadvertising authentic content and that cost is fixed, when you impose a tax on advertising, probably going to change the trade-off between advertising return versus the authentic contents return, although both may affect following behavior. So I guess my guess is in that context, the advertising tax may not be neutral either. So that's just my hunch.

The second comment is on the FTC advertising disclosure guidance. So I agree with Professor that the FTC's action going to affect the return to disclosed ads, as well as return to nondisclosed ads. I think the paper treat those two as two free parameters, and in reality, these two are actually linked because of consumers' belief, right? When you allow some to be disclosed, it's going to change people's perception of what is really behind the nondisclosed ones? So in that sense, the two tools probably are linked. I think it will be interesting

100

to explore the connection between those two.

Another thing I want to emphasize is that in the basic model, we sort of assume, okay, here's a fixed pie, we're just talking about how to divide that time. While each party may get zero or a positive fraction of this one, but in reality, the pie that's available for the influencer to grab is actually bigger than the real pie. You could sort of pedal up your advertising so that the followers sort of will pay higher price to the advertiser, who will kickback you a higher fraction of advertising revenue, but that product turns out to be much worse than what you advertised, so you sort of grab an inflated pie, and leaving a negative part to the follower, and that is not allowed in the basic model, but this inflation from the real pie is something I think really worry by policymakers, because your action in advertising generates damage to the followers, not just in the sense that they do not receive good advice, but also sort of generated damage negatively and impact them in terms of higher price or other forms. So I think that is a -- will be interesting extension. My hunch is that that is more than just changing outside option, because this affect the influencer's payoff directly.

Okay, about opt-in disclosure, the

recommendation is that FTC only enforce disclosure for small influencers. The big influencers, they will choose nondisclose, and they will be sort of let be in the market, and they -- and I understand the economic logic there, because the top influencers are in the harvest mode, and harvest is kind of the motivation for them to sow good advice beforehand, okay? So I understand the hunch that we need to keep the incentive there in order to generate good advice.

However, this is very much against the practice I have seen at FTC. For example, FTC has caught Kim Kardashian in the Skechers case, where Kim Kardashian has been involved in some deceptive advertising for Skechers shoes. FTC also send out warning letters to 21 social media influencers in April 2017. I understand there is a new round of warning letters going out just recently. So that is targeting big influencer rather than small influencer. So this is sort of quite opposite from the opt-in disclosure recommendation, and that's -- at least on the policy ground, it was justified by the potential large damage to sort of the negative return to the followers I talked before, and I think intuitively, that could be better for a big influencer, because they have a lot of followers today. So I think it will be good to

reconcile these two, the model recommendation versus -- versus the FTC practice.

And lastly, Professor Mitchell has not talked about search engine bias, but in the paper, there is a lot of models try to talk about search engine bias. I will just talk briefly. The paper models search engine bias in two ways. One is that the market power would increase the higher -- would mean a higher payoff in advertising revenue, which I agree, but it also assume the market power would imply a higher sort of value of good advice, and that's something I'm not sure I follow. So it might be good to sort of justify that.

Another way to model is sort of assuming there's additional income coming to the influencer, independent of the advertising behavior, which is modeled as a constant added to the income to the influencer. My question is, I would actually even want to think of this V, the constant return to the influencer, to be something that affect the follower's behavior. So I am thinking the V as kind of the value it can generate by authentic contents, which we have seen in a lot of examples of Twitter or search engine or other things. So that would have a big impact on the followers. We have seen a lot of arguments saying

that, okay, we need advertising revenue because that support us to create authentic contents, which generate a lot of good value to the followers, which sort of encourage the following. So I think it would be good to sort of bring the two together and allow both to affect the following behavior.

So overall, I think this is really a novel and general paper that applies to both antitrust and consumer protection issues. It has a lot of interesting insights. I've listed a few of these, but we encourage you to the read paper. It is a really fun intellectual exercise, and I hope the future version would get closer to the real business model and FTC practice. I know the Professor in going in that direction, so I really look forward to seeing the update.

Thank you.

(Applause.)

MR. WILSON: Thanks very much. We have got time for just one or two questions before our break, if anyone has one or two.

Oh, sorry, Jonathan.

MR. ZINMAN: Matthew, I think there's some evidence -- I'm thinking of some papers by George Loewenstein and coauthors -- that under the type of

disclosure regime you have in mind, that some consumers can end up being excessively trusting of the -- of the sender of -- of the provider of the advice. So I'm wondering if you think that could be materially impactful in your setting, for example, whether it would move up the optimal time of harvest and what implications that might have.

MR. MITCHELL: You know, I mostly did most of my comparative statics on the influencer side. On many things, there's a sort of almost equivalence. That is, you know, something that -- like you're thinking, that makes the total pie shrink or grow in a different way. You're thinking it also affects, like, the slope of the division between the two, because if you're overly trusting, that -- you know, that affects the division between the two. So it's probably a lot like -- I haven't exactly done that explicitly, but I think it's a lot like those things.

I want to stress, like, in the -- in the -- I don't really have a way to behaviorally think about exactly the words "too trusting," except that I do have a way to think about the possibility that they can't sort out one of type signal from the other, and that may make them respond excessively. Like the one I was thinking about was more that under disclosure

105

policy, you don't know when to follow advice that may be good advice but that has hashtag ad on it, but you could just as easily put in the reverse, and any cost like that of disclosure policy that's going to lower the total pie is going to in some sense -- I think going to have some of the same implications as M in the model.

MALE AUDIENCE MEMBER: In your model -- I mean, it's a moral hazard model where the agent cannot get any reward unless he shirks, right? I mean, that's when the agent gets a reward. So eventually the agent has to shirk. It's the only way the agent can be compensated and get some utility.

In -- in reality, I think that part of what these influencers have is -- I mean, they do have some value of being there and, you know, being influencers, their egos or the attention, the number of followers. There might be other ways in which they're compensated, by the fact that they are very influential, and not necessarily through ads that they need to, you know, steal from people.

I mean, I guess that -- and in your model would imply, you know, having, like, some flow utility that the influencer gets, what would be the consequences of that, and --

106

MR. MITCHELL: That one is literally an extension in the model, and so in the title slide, I said a theory of Kim Kardashian and Charlie sheen. The story there is that -- two things. So suppose that there was just a fixed benefit of having a follower, separate from the ads you could run. So I'm imagining -- like, I was thinking about the ego effect, maybe that you like having a lot of followers, and that's where I think of Charlie sheen.

So what does that do? Well, that's unambiguously good for followers because it makes it easier to get incentives because the threat of leaving them is even more severe. So that explains why attention seekers like Charlie sheen get attention on the internet, because they make good advisors in this model.

It is not unambiguously good for the influencer, though, because it makes it harder to extract through the shirking channel, because it's easier to get incentives on than to not shirk. So that kind of thing is unambiguously good for followers. One way to think about that is that Google could use as a sort of defense, that we need to -- we're good, because we want people coming to Google, and that makes us want to give them good advice in the

107

organic search results. So that would be sort of the policy way to think about your comment.

MR. WILSON: All right. Thank you very much. I think we are going to take a break now to try and stay on schedule. Let's reconvene in just a little over ten minutes at 11:35. Thanks very much.

(A brief recess was taken.)

108

KEYNOTE ADDRESS

MR. ROSENBAUM: We are going to get started, if everyone could please be seated.

So our first keynote address is going to be given my Professor Jonathan Zinman, who's a Professor of Economics at Dartmouth College, an academic lead for the Global Financial Inclusion Initiative of the Innovations for Poverty Action, and a co-founder of their U.S. Finance Initiative.

His research focuses on household finance and behavioral economics, and he has papers published on economics, finance, law, general interest science, and his work has been featured extensively in the popular and trade media as well. He applies his research by working with policymakers and practitioners around the globe, and it's our privilege to have him here to serve on the scientific committee and to hear his keynote address on "Modeling With Behavioral Consumers: New Evidence, New Tools."

MR. ZINMAN: Thanks. Thank you for having me at your conference. I very much -- given my fields and my interests, I very much feel like a guest here, which is quite exciting. Lots of acknowledgments, but I want to especially acknowledge the FTC crew, Ted, Nathan, and Daniel Wood, for helping me think through what might be interesting to you all, to this audience. I hold them harmless, however. If you find this talk boring, that's on me, not on them, definitely.

I also, of course, want to thank many coauthors who have provided many, if not most, of the inputs that this talk is based on, especially my coauthor Victor Stango, who is my co-conspirator in much of the work that I'm going to be talking about today.

So, okay, all right, game plan for today. So I'm going to be talking some about a big new project with Victor Stango and Joanne Yoong, which has produced two papers -- two working papers so far, with many more to come hopefully. I want to tackle two broad questions that hopefully I can convince you are interesting and worth considering.

One is why it's important to take behavioral biases in consumer decision-making seriously, all right, and I will at least briefly deal with a lot of the concerns and critiques about whether we should -- do behavioral factors actually matter out there in the wild when we have the types of repeat play and high stakes that we heard about this morning, for example? And if we are to take behavioral biases seriously, how do we do so from a modeling perspective?

So one -- one example of this would be, well, what should the behavioral, in a behavioral I/O model, look like? Hopefully that's an interesting question to contemplate for at least some of you.

All right. So to get started with some motivation -- all right, let's say we want to design or evaluate a policy, all right? So before we get into something that's close to my heart, we could also be thinking about designing or evaluating a consumer protection policy for one of FTC's markets, the influencer market or the used car market or eBay, all things we've heard about this morning.

All right, closer to my heart and my work, let's say we want to evaluate the CFPB's newly issued, as of four weeks ago, final rule on the very controversial payday loan market, or better yet, let's back up and model and conduct welfare analysis to diagnose whether and how we should be intervening in the first place.

All right, so when we're doing this, we need to decide whether we should consider behavioral factors, and that might influence consumer decision-making in our model, in our model of consumer behavior, in our model of how suppliers are going to respond given how consumers decide, in our model of how policy is going

to influence both types of parties.

All right, and one important question that I will largely punt on today in the interests of time and also statistical power is which behavioral factors to consider. So I'll -- this will come up again, but just to start fixing ideas, one of the challenges in behavioral economics and in applying behavioral economics is that there is a potpourri or panoply of biases that are thought to potentially substantially impact consumer decision-making, everything from present bias discounting to many varieties of overoptimism to loss aversion, to exponential growth bias, to statistical biases, like gambler's fallacies, and so on and so forth, all right?

So one of the things, without directly answering this question of which biases matter in which context, I'm going to talk about measurement tools and methodological approaches that can help us deal with this flowering, deal with this proliferation.

Okay, but first, let me answer the threshold question so that I can hopefully hold your -- continue to hold your attention for the next 20 minutes or so, which is what -- at a high level, what's the evidence on whether this stuff actually matters out there in

the wild? All right, and let me -- and so I'm starting by addressing any skeptics.

All right. So, first of all, there is evidence -- still not enough in my view, I go into this pretty -- and went into this project I am going to be telling you about today pretty militantly agnostic, particularly by the standards of practicing behavioral economists, but let's just say there is mounting evidence that behavioral tendencies, tendencies towards bias in consumer decision-making, at least, these tendencies are closer to ubiquitous than anomalous, and we have some new evidence on this, and we are standing on the backs of, among others, two recent Nobel Prize winners.

All right. There's also evidence -- again, not enough for my liking, again, one of the reasons why we undertook the project I am going to be telling you about today -- there is also evidence that the influence of behavioral factors on consumer decision-making do not disappear as stakes rise. There is actually ample evidence from the field -- from field settings at this point that they do influence large stakes decisions.

All right. Perhaps most shockingly, there is a fascinating and relatively new theory literature, not

113

yet, to my knowledge, really brought to the data, but there's a fascinating new theory literature exploring how and why consumers do not necessarily learn to debias themselves. They do not necessarily learn about their biases or how to correct them.

And one of the reasons I was excited about the panel this morning is it gets us thinking about delegation, all right? The panel this morning illustrates that delegation, intermediation, intermediaries who are providing information, maybe misinformation, persuasion, this is all -- it's all nontrivial to understand how this affects market outcomes even if we assume classical -- classically rational consumers. Imagine allowing for behavioral tendencies among consumer decision-making, all right? So that's a long way of saying we really don't know -- and there's actually some empirical evidence suggesting that we should be skeptical, but let's be more agnostic -- we really don't know whether delegation and intermediation serves to functionally debias consumers and cure the would-be impacts of behavioral biases on decision-making.

Okay. And the last bit of motivation for true believers, even if you are already convinced that behavioral biases influence consumer decision-making,

114

we still need to do behavioral I/O modeling, certainly when we want to understand the impacts of potential policy interventions, because evidence is mounting, both theoretical and empirical, that seemingly intuitive treatments, seemingly intuitive interventions can actually make things worse, particularly when there's limited enforcement.

Okay. So the broader motivation here is, you know, apart from any particular market, whether we're focused on payday lending or used cars or whatever, the broader motivation here is developing tools and evidence to inform how we should use those tools about how we can build portable models that reasonably and usefully capture behavioral consumers, all right? So I'm going to be -- I'm going to be talking today a bit about different approaches to specifying -- designing and specifying models, and what we're going for here is building more workhorse, portable behavioral models, okay? So that will be -- this is going to be my last four slides in approximately our next three or four papers, hopefully, which are going to be summarized at a high level in these last four slides.

But first, I want to introduce this project that Victor and Joe Ann and I have been working on for years and are -- and that is finally bearing fruit in

115

the form of some working papers. So what we do in this project, the Multiple Behavioral Factors Project, is we collect data on over a thousand representative U.S. consumers using the RAND's American Life Panel, and we're -- in this rich data set, we're collecting data on various behavioral decision-making tendencies of these consumers.

So rather than doing what behavioral economists typically do in lab-type studies, which is bring someone into the lab and hammer away at measuring one particular bias for, say, 30 to 60 minutes, with a very repetitive set of tasks in a lab, and so rather than just try to measure whether people exhibit time consistent discounting and, if not, whether they're present-biased or future-biased, we're going to do streamlined versions of that and measure 16 other potentially behavioral influences on decision-making.

So in addition to measuring discounting and any discounting biases, we're also going to try to measure loss diversion; we're also going to try to measure exponential growth bias; we're also going to try to measure statistical biases; we're also going to try to measure limited perspective memory; we're also going to try to measure three different varieties of overconfidence; and so on and so forth.

116

All right, we do this, as I've already intimated, using what behavioral commits refer to as direct elicitation. So for the uninitiated, what's direct elicitation? It's putting people through stylized tasks that are meant to reveal their decision-making tendencies. The analogy here -- and this is -- and I should emphasize, you know, as with any methodology or as with any measurement technology, direct elicitation certainly has its pluses and minuses. We certainly think of it as a strong complement to various other methods of measuring or inferring influences of behavioral factors on consumer decisions and market outcomes, but just by way of sort of motivation and history, there's -- there's an analogy here to a much longer history in the social sciences of intelligence testing and personality testing, all right?

So one can try to infer someone's intelligence or cognitive skills by looking at things they do out there in the wild, right? So you could try to infer cognitive skills from how people perform on their job, for example. Well, it turns out you can also try to infer and measure cognitive skills and learn a lot about people by putting them through stylized tasks or tests, all right? So we're on the stylized tasks and tests side of things here.

Okay. And so along with collecting this rich -- rich in a broad sense, not rich in a deep sense -- but in tandem with collecting data on these 17 different hypothesized behavioral influences on decision-making, we also collect a lot of other information on people taking our surveys. Specifically, we try to measure what one might think of as standard or classical decision inputs or factors; cognitive skills, for example. You know, we do -- we implement some standard short versions of intelligence tests. We also elicit classical measures of preferences, right, so patience, classical risk attitudes. And, of course, we also have a lot of demographic information on these folks, including things that would be important in, say, any life cycle model of consumption and consumption savings decisions, all right?

The great thing about this survey and the panel we're part of in this survey is you also get a lot of rich data on decisions people are making in their real lives, assuming they're reporting reasonably truthfully, and we worry a lot about that. Being household finance people, in our modules, we're particularly focused on household finance, but there

are and will be in future iterations of our working papers many other outcome domains that one could look at here, human capital type stuff, health type stuff, et cetera, et cetera.

Okay. And so what's coming out of this project and our working papers and future working papers are sort of two classes of things. One is new tools for measuring behavioral influences on decision-making. One of our -- one of our papers that's done is partly focused on showing that these streamlined elicitation methods that we use to measure 17 things instead of one thing that might be behavioral influences on decision-making, so part of what we do is demonstrate in various ways that these streamlined elicitations actually do produce useful data, all right?

So what we have now is a suite of low-cost, direct elicitation tools that are portable to a broad variety of data collection settings. You know, part of what we end up arguing here is you no longer need to bring people into a lab and do extensive, expensive, high-touch elicitations to learn useful things about how behavioral tendencies might be influencing consumers' decisions. You can use our streamlined elicitations instead.

We're very worried about measurement error in

all aspects of our data. I'm definitely on record and published about being worried about such things in prior publications on -- in terms of survey data. And so a lot of what we're doing is working on developing new -- or at least new for economics -- types of measurement error corrections and also comparing them to more standard and well understood measurement error correction techniques.

We construct new summary statistics for -- at the consumer level for capturing behavioral decision-making tendencies. I'll talk in a couple slides about how these end up being useful. And so -- and along with the new tools, of course, we also have some new evidence on what we think are some foundational and still largely open empirical questions. So you can use our data to look at the prevalence and heterogeneity across consumers of these 17 different behavioral factors.

It turns out many of these factors are quite prevalent. They are also quite heterogenous across people. Being behavioral on one dimension, particularly in directions that have been the focus of prior literature -- so, for example, being present-biased instead of future-biased, having a preference for certainty instead of a preference for

uncertainty -- underestimating the power of the large numbers as opposed to overestimating it.

It turns out that if you're behavioral on one dimension, you're substantially more likely to be behavioral on other dimensions. So I'll talk towards the end about some possible implications of that finding.

These -- these measures -- these measures of behavioral stuff turn out to be statistically as well as conceptually distinct from classical factors, both in terms of measures of fit and measures of conditional correlation with the types of outcomes we might care about. And as just alluded to, many of these behavioral biases do turn out to be correlated with real-world decisions and outcomes, like, for example, various measures of household financial condition, and that's conditional on our measures of classical factors, demographics, everything else we observe about these folks.

Okay. So what do we -- what do we do with this? How can we model behavioral consumers? How can we capture something useful about behavioral tendencies in decision-making, understanding that this generates substantial additional complications if we're trying to build an equilibrium model that allows for supplier responses, that allows for treatment effects of policies, or other interventions. So what should we do with all this?

Well, one approach -- and far and away the standard and most popular approach historically, whether in behavioral I/O or behavioral anything, in economics -- is what has been referred to -- and not charitably -- as the silo approach, right? So that's -- you know, there are dozens, maybe even a hundred at this point, of behavioral biases that have been hypothesized and in some settings suggestively shown maybe to influence or at least correlate with decision-making. The approach so far mostly has been, well, we're just going to deal with these one bias at a time.

All right. There are a lot of folks who, quite understandably, are concerned about this, right? It is not very congruent with building portable workhorse models of behavioral influences on decision-making. Drew Fudenberg maybe has the most, I think, high-profile and incisive critique of the hundred biases/hundred different models problem.

But this is a valid way of doing business if behavioral biases are separable from each other in terms of how they influence consumer decisions.

Okay, so going back again to something that's close to my heart, in some of my other work, say I have reason to believe that that overoptimism about repayment -- and, you know, there might be some behavioral stuff underlying that forecast error, which we could talk about later -- but let's say in a reduced form way, overoptimism about repayment is an important feature of payday loan borrower decision-making that I'm worried about as a policymaker.

Can I just model that and ignore any influence of present bias, ignore any biases that might result from people with present bias discounting getting tempted by quick cash when they drive by one of the countless payday loan storefronts or when they encounter one of the countless ads or links to online payday lenders? Can I ignore that safely?

Well, until now, there's been very little evidence to guide us on this modeling decision. With our data, you can begin to tackle this empirical question, and the evidence we're finding thus far is quite encouraging, surprisingly so, actually. So basically if what you do is you start by estimating richly conditional correlations between, say, some outcome or condition you're interested in, say an

index of overall household financial condition, which is basically capturing some sort of -- you know, a series of correlated signals of wealth or financial stability, all right?

So you start by estimating correlations between that outcome index and single behavioral biases, okay? When you do that, we find patterns of correlations that line up very nicely with standard behavioral silo theories. You know, present bias guys look like they have worse financial conditions, conditional on everything else, right? Guys with limited memory, per our stylized tasks, have worse financial condition, conditional on everything else.

If you then add a vector capturing everything else we observe about these folks behaviorally speaking from our elicitations, these results do not change at all. All right, I did some fun effects there, because I think this is a potentially profound and exciting result, all right? It basically supports standard operating practice in most of behavioral economics.

It suggests that, at least in the one outcome domain we've looked at so far, and subject to all the caveats of -- about correlational reduced form analysis, it suggests that behavioral biases may,

indeed, be separable in ways that are amenable to siloed modeling where the silo -- where the silo, of course, you know, may accommodate two or three biases that interact, all right, but it's -- you know, the siloed approach is basically one or few biases at a time, not a dozen or a hundred at a time.

Okay, all right. There's another approach, which is to say, well, consumers are behavioral; we're not sure how or why. All right? This is the reduced form behavioral sufficient statistic approach, all right? So in these models, there's a wedge between decision utility, what people think their utility is going to be when they make a decision, and experience utility, what actually ends up happening, all right?

Reduced form models often get a bad rap in economics, but as I hope to show you, these models can be very useful, and in other -- and other fields are very happy to make and explore distinctions between emergent versus fundamental models, right? And so this is an emergent model. This is a model where we have a core specification of how people go awry due to behavioral influences on decision-making without modeling all the fundamentals of exactly how they're going awry.

So how do you do this? Well, fortunately, for

125

all of us, Raj Chetty and my new coauthors, Hunt Allcott and Dmitry Taubinski, have some great papers where they -- where they develop and explain this tool kit far better than I could in two minutes or less. I have not yet, by the way, seen this approach deployed in behavioral I/O, although it's possible I've just missed some interesting papers.

But anyway, using this reduced form approach, people are behavioral in some way. We're going to specify that coarsely, in reduced form. Even this approach relies on some key assumptions. These key assumptions have also have not been validated or invalidated empirically. Again, you can take the data and Victor and Joanne and I have generated and poke at these assumptions, all right?

Again, the findings are encouraging for the most part, although not -- although not universally in the case of the reduced-form, sufficient statistic models. So one key thing you need for these models to work and for them to make sense is you need to posit it within consumer correlation amongst different behavioral biases. As I said, on the last slide we had that or two slides ago we had that.

For -- we -- we take that as a jumping-off point and then actually construct simple

126

consumer-level summary statistics, aggregating across behavioral biases within -- within consumer. In doing that, you find support for another key assumption these models have, which is that people actually need to be biased, right? And to my mind, this is actually what behavioral economics is all about. It's not about people making mean zero errors. It's about people tending to make errors in a particular direction, exhibiting bias.

So we find that, and you can use our summary statistics to illustrate that. Moreover, these summary statistics end up being strongly conditionally correlated with outcomes, with outcomes and decisions in the field.

All right. The one caveat here -- and I think, to my understanding, what really complicates trying to use these models for policy applications -- is that when you have heterogeneity in how behavioral consumers are, it's actually quite difficult, quite a heavy lift to identify the average marginal bias distribution you need to do welfare analysis, all right? So you really -- to make good use of this method, you really need to have good data and good identification that allows you to sort of walk down the behavioral demand curve, and that can be

127

challenging, although Hunt and Dmitry do a very clever and thought-provoking job of this in their AR paper on the light bulb market.

All right. So a third approach, which is very much still under construction, is grand unification, all right? So is there something fundamental about human decision-making that produces these 17 or these hundred different behavioral biases and their links to decisions in the real world? It's not crazy to think this could be the case. I mean, we could draw inspiration from other fields as far-flung as physics, but closer to home, this is what -- this is very much what social scientists in related fields on decision-making have been discovering over the last many decades.

We started over 100 years ago with the model where there were basically countless cognitive skills and ways people could be smart or skilled. That has been distilled to what's sometimes referred to as the G factor, smarts, intelligence, general intelligence. Similarly, in personality psychology, all right?

We find some encouraging results, one of which I have -- I have already mentioned. Taking it a step further, if you subject our data on multiple behavioral biases to factor analysis, it does look

128

like there is a single common factor underlying the 17. So that's very exciting and would seem to bode well for prospects for a grand unification, but so far we're finding the glass is half empty in the sense that that common factor does not seem to help us understand real-world decision-making or outcomes conditional on what we already observe about people, but there's still much more work to be done on that margin.

Okay, so last slide. Summing up what to make of all this and how some of you might be able to think about using this evidence and these tools going forward, so you have a setting, you have a market you're interested in, where you or the policy folks you're working with have priors about a behavioral bias or a set of behavioral biases that affect consumer decisions and possibly welfare. What can you do?

Well, you can use our tools to cheaply and directly measure the behavioral biases of interest in the market you're interested in, to see whether they're prevalent, to see how much heterogeneity there might be. You can then use that data, the data on the empirical distribution of your bias or biases of interest, and data on statistical relationship between

129

that bias or those biases and the outcomes you care about and the market you care about to inform your modeling decisions about how to model competition equilibrium policy impacts in this market you care about, right?

You can use this data to inform whether you should or could build a behavioral silo model, where you're just focused on one bias or, say, the interaction between two biases, or if it seems like there may be many biases in play, which are positively correlated within people, and so on and so forth, you might want to go the reduced form behavioral sufficient statistic route.

All right. Eventually, hopefully, we or one of the other teams working on the grand unification question will have a third option to offer, but I think we're some years off from that. And I would say in terms of the overall approach on this slide, Hunt, Dmitry, and I are putting our money where my mouth is today and trying to use just this approach in various markets at this point, and I hope others will join us on this journey.

Thanks.

(Applause.)

MR. ROSENBAUM: We have time for about one or

130

two questions. Okay.

MALE AUDIENCE MEMBER: Regarding welfare, how do you evaluate -- I mean, what is the welfare criterion that you would use? I mean, these people -- I think in some ways you might be regarding them as they have -- they don't know how to make decisions, and so there is some utility function that they really have, but yet they're behaving in a way that is inconsistent with that, or be more agnostic to what are really the preferences. Maybe they have preferences over, you know, decisions, actions, and so how do we even go about thinking about welfare?

MR. ZINMAN: So for some behavioral biases, this is -- the answer to that question is relatively straightforward. So there is a distinction between biases and preferences, which raise the thorny issues that you just mentioned, and biases in beliefs or in the processing of information, right? So for the latter, it's relatively straightforward. It's -- you know, it's usually reasonable to use the unbiased benchmark for our welfare analysis.

When people have behavioral preferences, it is -- it is far thornier to deal with. The most -- you know, I think the -- in recent years, the greatest focus in behavioral economics has been on

131

time-inconsistent discounting, in particular, which is actually sort of a mishmash of preferences and beliefs if we really want to get per in this case at this about it, but anyway, the -- you know, sort of the standard operating practice, to the extent there is one, is to -- is to imagine that the behavioral guise would, in fact, be time consistent and would prefer to be time consistent.

MR. ROSENBAUM: One more.

MALE AUDIENCE MEMBER: Hi. Just a question on the summary, these three approaches, the silos versus this grand unification. If I'm understanding it right, it seems like if the -- if the silos work, then grand unification can't work, because what the silos are depending on is the fact that the part with, you know, behavioral bias A that's correlated with behavioral bias B doesn't explain the outcome variable of interest, that it's the common component that's uncorrelated with the outcome, and grand unification requires that all these, you know, 17 or 100 bases, there's a component of them that together is correlated with the outcome. So how can silos work and still there be hope for grand unification?

MR. ZINMAN: So I -- I suspect -- I suspect you are right, that if one works, the other doesn't. I

132

would hedge in two -- in at least two ways, though. One is we haven't proved that. We haven't fully worked that out, nor has anyone else, and there could be some subtlety and nuance that makes this worth seeing whether one could prove it.

The second thing is that, you know, all of the evidence I just presented to you from our stuff, at least, including the evidence of validating the silo approach, is new and preliminary and consequently should be taken with a grain of salt.

MR. ROSENBAUM: Thank you very much. Let's thank Jonathan.

(Applause.)

MR. ROSENBAUM: Now we are going to break for lunch. There's food outside, and we'll take 25 minutes for lunch. Let's try to be back at 12:45 for our next panel.

(Whereupon, at 12:22 p.m., a lunch recess was taken.)

AFTERNOON SESSION
(12:48 p.m.)
PANEL DISCUSSION

MR. ROSENBAUM: All right, if everyone could be seated. I'm going to turn the microphone over to my colleague, Keith Brand, who's going to introduce the next panel.

MR. BRAND: Good afternoon. Welcome, everyone. My name is Keith Brand. I'm an economist with the Federal Trade Commission, and I will be chairing our panel discussion this afternoon on cross market provider mergers.

As many of you are likely aware, several recent empirical and theoretical studies examined the price effects of cross market mergers between healthcare providers. For the most part, these studies consider whether mergers between healthcare providers in nonproximal geographies lead to higher prices even though the providers are not close substitutes for patients at the point of service.

I think it is fair to say that the empirical analyses and the literature do provide credible evidence that prices have increased following such mergers, and while the literature has explored several mechanisms that could explain the empirical results,

it is perhaps less clear that we have good evidence on what mechanisms are likely to be the most relevant.

We have assembled an outstanding panel this afternoon to discuss the literature on cross market mergers, what research has been done, what we think we've learned so far, and what are the most likely important next steps in the literature.

First, to my far left, we have Marty Gaynor. Marty is the E.J. Barone University Professor of Economics and Public Policy at Carnegie Mellon University and the former Director of the Bureau of Economics at the Federal Trade Commission. He's also a founder and a former chair of the Governing Board at the Healthcare Cost Institute.

Next to him we have Matthew Schmitt, who is an Associate Professor of Strategy at the UCLA Anderson School of Management.

Next we have Greg Vistnes, who's a vice president at Charles River Associates. He has also served as the Deputy Director for Antitrust in the Bureau of Economics at the Federal Trade Commission and as the Assistant Chief of the Economic Analysis Group at the Department of Justice's Antitrust Division.

Finally, we have Matthew Lewis, who's an

Associate Professor of Economics in the Department of Economics at Clemson University.

So we have organized our discussion as follows: First, each of the panelists will provide some opening remarks on the topic, and then we've grouped together four topics for discussion after the opening remarks, and we plan to leave about 15 minutes or so for questions and answers at the end of the panel.

So we'll start with Greg Vistnes.

MR. VISTNES: Okay. Well, thank you very much for the opportunity to be here and speak here today. I think this is a really important topic. I think, given all the interest that's out there, both among economists as well as some of the different enforcement agencies, it's a very ripe topic to have this sort of a discussion.

I just want to sort of open up with what I think are three of, to me, the most important issues about some of these cross mergers and the enforcement issues. First of all, why are we looking at it? What is it that makes this, at least to many of us, such an interesting topic?

Secondly, do we have a theory for any of these concerns? And maybe even, why do we really care if there's a theory? Is that important or not?

And then third and related to it is, what are some of the policy implications about pursuing an enforcement agenda? They are all sort of wrapped up with each other.

So, really quickly, why are we looking at it? Well, to me at least, we're looking at it because we've heard complaints -- I've heard complaints -- for over a decade from managed care plans saying these things are bad; these hospital systems with hospitals even in different markets, they just make us all in a worse situation. And there's never really been a good economic theory to explain that, but then part of what's recently come out from both the Matts on either side of me and from others as well is now there's some empirical evidence to back up those concerns, that what people are saying, there seems to actually be some truth to it, that some of these hospital prices for chains seem to be higher, may be due to this -- call it a cross market effect, but we still don't have a good theory. What the heck is the theory?

The theory that we're looking at is not the traditional vertical theory. It's not foreclosure, it's not bundling, it's not tying. It's something different. Well, what is it? You know, here some of my biases are probably starting to come out, but it's kind of like the conglomerate effects theories of the 1960s. It's kind of like the portfolio power theories of the 2000s that Europe pursued for a while. It's not really clear what the heck this is. So what are we going to make of it?

Well, there seem to be at least two aspects of the theory, and we're going to talk more about the details of the theory, but there seem to be sort of two parts of, if there's a theory to explain cross market mergers that we've come up with so far, that somehow the theory has to explain linkages across these markets, and the linkage is not coming, by definition, from patient flows like it is in the traditional, but there has to be a linkage to make cross market effects work.

And then secondly, it has to be a really special kind of linkage. It has to be -- and, again, we will get into this in gruesome detail -- it has to be concavity of a linkage effect, concavity of profits or superadditivity.

But then it turns to the other thing is, you know, yeah, to heck with it. We have empirical evidence that the effect is there. We have got complainants. Why do you need a doggone theory with these economists concerned about proving what everyone

knows is true? Well, there are a lot of good reasons for it, and we'll, again, get into hopefully a lot of it in this discussion, but really important, at least to me, is we need to be able to offer guidance. What, in essence, we have now is I'll say we have a possibility theorem. We have proven that it is theoretically possible. Is it likely? Where is it likely? What's the magnitude of the effect?

And importantly, from the providers' perspective, what the heck can they merge with if they don't know which ones are going to be challenged or not? Some sort of guidance has to be provided there. So how can we give them that kind of guidance?

And then the last thing that I want to mention that I think is, again, super important with policy implications are, what are the limiting principles? Where do we stop? Is it just cross market with respect to hospitals in different geographic markets, or do we especially start looking at product markets, because the theories will probably extend pretty easily.

Do we start caring about acute care hospitals and children's hospitals and psychiatric hospitals getting it together? What about acute care or what about inpatient versus outpatient? What about

hospital physician mergers? What about multispecialty clinics?

And then, heck, why stop here in healthcare? We have got the world to explore. We have got cable TV. We've got all sorts of markets where we can apply this theory, where is the principal issue payer complaint? I hope not. So we need to wrap this through.

And then why I think this is such a critically important issue or topic for discussion is we have got a bunch of really bright economists here. I don't think the theory is out there yet. There's a lot of reason to think there may be a concern, but if anyone can figure out whether or not to accept or reject a theory, I think that's a great research opportunity that's going to have some real value for folks.

MR. BRAND: Okay. Thanks, Greg. Let's next turn to Matthew Lewis.

MR. LEWIS: Okay. So I'd like to -- just given that introduction, I think I'll spend my time just giving some background on the recent empirical evidence by going over the results of my two papers, and then I'll leave it to Matt to discuss a few of the others.

Actually, the -- I have written two papers,

both with Kevin Pflum, on this topic. One was more of a theoretical -- a structural paper which built off of the structural model of hospital MCO bargaining that's commonly used to study within-market mergers, but thinking about a new twist, which is the extent to which being a member of a hospital system might impact the bargaining power of the hospital, where bargaining -- bargaining power is -- when saying bargaining power, I'm referring to the Nash bargaining weight, which we will talk more about, but -- so that's distinct from any local sort of market position of the -- of the MCO and the hospital.

And what we find there is some evidence that hospitals in systems do have higher bargaining powers and that -- and that bargaining power is increasing in the size of the system, even if the system partners are outside the local market. So this is starting to -- that paper does not establish any causal effect of being in a system and how that impacts bargaining power, but it's suggestive that maybe there's this opportunity to link up with hospitals in other markets and somehow increase my negotiating ability through this bargaining power parameter and get higher prices.

And so that inspired the second paper that we have, which went on and specifically looked at observed cross market mergers, so over 100 of these mergers, examining what happens when a stand-alone hospital is acquired by an out-of-market system that has no other partners in the local market, and what happens to the prices of that stand-alone hospital, the acquired hospital, what happens to the prices of their local rivals in that market.

And we show that, on average, the prices of those acquired hospitals do go up by something like 17 percent, on average, and also you see an increase in the prices of their neighboring hospitals. So there's some suggestive evidence that -- again, that there's a -- basically a -- you know, some sort of softening of competition here in the sense that prices and profit -- price gross margins are going up here, and what -- and based on this evidence and some supplementary analysis, we argue that what it -- the patterns that we see in these price increases appear to be most consistent with the possibility that the bargaining power of these hospitals has changed with the merger. Basically, that they are somehow acquiring an increased ability to bargain -- bargaining sophistication, some increased ability to gather more of the rents available.

So I think several other papers have since been

put out -- I guess working papers, too -- that used a similar difference-in-difference approach to study cross market mergers. The -- each of these studies -- I think this is interesting. Each of these studies studies a somewhat different set of hospital -- of cross market mergers and looks at different firms when they're evaluating the price effects of those mergers, and so in that sense I think these studies are really complementary, and what I think we can do -- you know, so, for example, we focus on acquisitions of stand-alone hospitals, and we argue that the evidence there suggests maybe that those hospitals acquire a stronger bargaining power in that acquisition, but other types of mergers -- you know, the evidence from these other studies suggests that there may be evidence that some of these other mechanisms that Greg talked about -- or we will talk about -- that there is evidence that some of these other mechanisms may be generating price effect -- cross market price effects in other settings.

So I think there's a lot of opportunity now to bring the results of all those papers together and think carefully about when and where we might -- we think we will -- we will see price increase -- price increases after these mergers and also what we can

tell about what mechanisms may be responsible in different settings. I think we can do that based on the evidence we have.

MR. SCHMITT: I'll just continue to give you a description of some of the empirical evidence we have for cross market merger effects, evidence in addition to what you just heard from Matt. So, first, Leemore Dafny, Kate Ho, and Robin Lee have a paper in which they examine hospital system acquisitions of other hospital systems, and their focus is on the outlying hospitals of those systems, so hospitals that are more than a 30-minute drive away from the closest hospital belonging to the other system. The goal there is exactly to shut down direct patient substitution between the merging hospitals.

They find that prices increase post-merger for the outlying hospitals but only when the outlying hospital gains a system member in the same state. So when a hospital gains a system member from out of state, they find no evidence of price effects.

What might explain that, Dafny, Ho, and Lee note that, while there may not be any direct patient substitution between the merging hospitals that they examined, A, and the hospitals may contract with the same insurer -- they call that common insurers -- and

B, the hospitals may both be valued by the same employer, imagine if you employ people in the northern and southern suburbs of a city and you offer a single insurance plan to your employees, you may care about both of those hospitals. They call that common customers.

They demonstrate theoretically that both common insurers and common customers can generate price effects in standard bargaining models, and both common insurers and common customers are more likely to occur in state than out of state.

Second, let me touch on my own work in this area. As regional and national hospital systems have expanded, they now overlap with one another in an increasing number of hospital markets. To give you just one suggestive statistic, about half of U.S. hospitals now belong to a system that operates in multiple hospital referral regions, which is a big market definition, and about a third belong to systems that have a presence in multiple states. So, in short, hospital systems compete with one another in multiple markets simultaneously.

In the literature, that's often referred to as multimarket contact, and there's a large body of theoretical work and some empirical work demonstrating

145

that multimarket contact can soften price competition. I have a paper in which I examine whether escalating multimarket contact between hospital systems, which has largely been generated by acquisitions without direct horizontal overlap, has led to higher hospital prices.

I don't know if it's productive to get into the details of the measurement there, but I find evidence suggesting that, indeed, more multimarket contact leads to higher hospital prices. In line with, I think, what Matt raised in his closing, in my view, what remains elusive is more direct evidence about what the true underlying mechanisms are. I think there are a few clear obstacles to really nailing down specific obstacles -- specific mechanisms empirically, but I'll stop for now because I imagine that's something we'll get into.

MR. GAYNOR: Great. Well, thanks.

So let me talk about some conceptual or policy issues at a high level, and then I'm sure we will get back to specifics. One thing that I think should be emphasized is that the issues that are raised here are not specific to healthcare. They are potentially quite broad and could apply in a whole bunch of other industries, lots of retail outlets, online outlets.

146

For example, if the manufacturer of Skippy Peanut Butter and the manufacturer of Charmin Toilet Paper wanted to merge, would that be a merger that would potentially be harmful to competition and worthy of the agency's attention? Mike Vita is looking at me like his head is about to explode.

We would -- under -- you know, under sort of consumer substitution, it clearly would not meet that criteria. If you give your kid a toilet paper sandwich for lunch, they'll like it even less than a peanut butter sandwich, but perhaps that's not -- that may not be the correct lens through which to view this.

So -- but coming back to healthcare, what we have at this juncture is we have fact patterns. Market participants say things that are consistent with cross market mergers, perhaps enhancing market power and harming competition. There are stories one hears from payers, in particular -- who, after all, are the people paying for this stuff -- and sometimes from health systems themselves.

Actually, could I get the slide, please, if I may?

So as folks may know, UNC Healthcare in Chapel Hill and Carolinas Healthcare in Charlotte, about 130

147

miles apart, about a two-hour-plus-ish drive on interstates, are talking about merging, and the CEOs are in the picture of the two health systems, and in an interview at newspaper offices, the executives said the partnership would give them leverage to negotiate better deals with insurers, at which point their lawyers' heads exploded.

So when -- what -- what does this mean about this merger? I don't know. I'm not opining on this merger. Obviously, it could be a beneficial or a benign merger or go the other way. The point is that at least the CEOs of these two merging entities who, arguably, very well may not be in the same geographic market -- you can take the slide down if you like -- seem to think that this is going to enhance their negotiating leverage.

Now, being CEOs and not Ph.D. economists, they didn't specify exactly whether that was due to concavity functions or shifts in relative bargaining rates. I don't know why. Somebody needs to do a better job in MBA strategy classes, I think, but anyhow -- and then we have the empirical patterns that Matt and Matt have ably described. So we see these things very carefully done, very, very competent, good research, where there are these fact patterns emerging

148

from the data.

As Greg said, it's not entirely clear what to make of these things. Now, in some ways, I think that's a blessing, right? How does science advance? One way science advances is we turn up stuff and we look at it and we don't know what to make of it, and so then we have to go back to the drawing board and think a bit harder about what's going on.

Now, it's certainly possible -- and there are stories we can tell, and, again, the folks who have been working in this research area do have some pretty compelling stories that rationalize the observed empirical patterns into some existing models, saying, you know, you just have to think about who the buyer is, and that makes a lot of sense, but I think that we're still not quite there yet. In particular, in being able to draw clear inferences about whether there's harm to competition and what the appropriate enforcement policy is.

So I think that we do need some further thinking about the underlying theoretical framework, and obviously some of that's technical, but really the question is what kind of behaviors are there that would generate this and then some tests that can sharply distinguish those behaviors from other kinds of plausible behaviors.

And then I think that a fuller -- a fuller -- what's needed is a fuller model, both theoretically and empirically, and, in particular, one has to include insurers in that model. That's been a big challenge in healthcare because the data aren't generally available.

Kate Ho and Robin Lee, who were mentioned previously, are some of the few people that have done that kind of work, and they have gotten the data, but they have not just gotten the data, they have thought hard about what the economics are and been able to specify and estimate very careful econometric models to capture that.

So I think we need more about that as well in order to be able to make progress on this front, and then I think a couple other just thoughts on that. One, it can be hard for academics to get a hold of data if the dataholders aren't willing to part with it, but folks in enforcement agencies do have subpoena power if there is an important issue. And while I'm not -- I would certainly never suggest that the FTC or any agency use those powers lightly, but when there is an important matter and it's important to know these things, there can be data available that otherwise

might be hard to come by.

Then I think, as Greg mentioned, looking at product markets is important because the general notion is not specific to geographic markets. It's about product markets, and there are certainly other industries. I was not entirely facetious about the peanut butter/toilet paper example. There are other industries where if this has validity, it would apply potentially with real force as well.

MR. BRAND: Thank you all very much.

So I am going to turn to two topics that address two of the main mechanisms that the literature has explored as plausible explanations for the empirical results. The first is, as Greg mentioned, the concavity or convexity -- as the case may be -- of insurance profits with respect to the providers included in its network and what may be driving that concavity or convexity. And the second, as Matt described it, as potentially the merger induces a shift in the Nash bargaining weight.

So I'm going to turn first to Greg on the concavity issue just to -- first to frame the issue, what we mean by concavity, how that connects with -- well, what you may think of it as a standard approach to analyzing healthcare mergers. And I know a number

of the panelists have some thoughts on is concavity or convexity more likely to obtain in the real world.

So, Greg, if you could kick us off on that.

MR. VISTNES: Yes. So I think there's a little bit of danger as folks up here at the panel right now are talking a little bit inside baseball, and everyone out there is saying, what the heck is really the issue you're talking about? So I might frame a bit the issue.

Standard merger analysis, you know, your typical widget merger, it's all based on the notion that there are substitutes, and the places where we care most about concerns are where one is a really good substitute, but what that really means is a consumer, when they're premerger, trying to decide between one or the other, they say, well, if I lose this one, I'm not that much hurt, because I can switch over to this other substitute, but if I lose the other one as well, because now I can't have either one, I'm a whole lot worse off.

So there's some concavity, or if you flip your graph upside, depending on what's on the other axis, convexity, but you have curvature. You have superadditivity in the sense, in a sense, that by losing the second one, I'm worse off. That's what, in

essence, we are trying to get at and what I think a lot of the theories in the hospital mergers is all about.

Now, if we're talking about two different hospitals in two completely separate geographic markets, where we're assuming by definition consumers don't go back and forth because they're separate geographic markets, different islands, if you lose one hospital, well, that's going to hurt everyone on that island, but why are they worse off if they lose the hospital on the second island? Why do we get that linkage? Why do we get that superadditivity or concavity in sort of the profit function? Why is it so much worse off? And so that's what a lot of the theory is all about.

I like to think of it a little bit as sort of the theory of holes, and from the managed care plan perspective, who's doing the purchasing and the contracting of all the hospitals, is, well, if they get a hole in one geographic market because they lose the hospital, is it that much worse if they incur a second hole in another geographic market? Are they getting increasingly worse off the more holes they have? And that sort of potentially opens the door to the theory having legs.

153

Now, a really important part -- and this is where concavity comes in, and I think it's also a potential danger when people are listening to managed care plans. It's really easy to ask a managed care plan, well, gee whiz, if you lose your hospitals on Island A, and then you lose your hospital on Island B, are you worse off? And they say, well, of course, you moron. How could we not be worse off? And people say, ah-ha, we've got it, cross market effects.

And then the economist really wants to say, well, gee whiz, what I really meant is, is it concave? Is there superadditivity? In which case the managed plan care plan again says, you stupid idiot, what do you mean?

So the notion that we're really trying to get at here is kind of the question you want to ask the managed care plan, is let's pretend you're negotiating with both of these hospitals at the same time. Now, you know you want both of them, and you know that if you lose either one, you're going to be kind of hurt, and now you're negotiating now with the hospital on Market B or on Island B, and then all of a sudden, someone comes in to your office, in to the negotiating room and whispers in your ear, hey, we just lost the hospital on the other island. Does that make you need

154

Hospital B even more?

If it makes you want to pay -- willing to pay them even more, if it affects your bargaining position on that other island because you lost the other one, then you have the linkage, and if you're willing to pay even more for it, you'll have concavity.

So the question is, how can we come up with a theory that establishes how this linkage and how the concavity can occur? And I'm not going to get into the details here. I can tell you that in playing around with trying to come up with a theory that is, I'll call it unbiased, that doesn't assume the answer, because it's really easy to come up with a theory of cross market effects where you basically implicitly -- and you kind of hide the fact -- but basically you're assuming this concavity -- but if you don't assume the concavity but have a really neutral market, it's really tough to get these effects. It's tough to get linkages.

And to get concavity? That's even tougher. And to get a theory that's unambiguously concave, as opposed to sometimes being convex, good luck with that. I haven't had any luck with that.

That leads us to the issue of, what is our theory going to tell us? What is it going to be good

155

for in terms of predictive? It goes back a little bit -- and I'll pass the buck in just a second -- is I think we do have the possibility theorem. We can show that. It's possible. But we don't yet -- I certainly haven't seen anything that gives much in the way of guidance about saying when it is or is not likely to be much of a problem, which, again, puts us back to the theory is having a hard time explaining what seems empirically, and from people's mouths, to be there. We've got lots of smoke, but we're trying to figure out, where the heck is the fire coming from?

MR. BRAND: Okay. Any other panelists want to weigh in on --

MR. GAYNOR: Yeah. So I think I'd perhaps be a little -- a little more positive, but -- but sort of -- one thing I could imagine doing is taking one of the stories that seems sensible on its face, and one of the stories that to me seems sensible on its face is you have got large regional or national employers, and they need to have these hospitals -- not just one, but both -- and then I don't think that writing down that model is terribly hard, but then -- then testing it empirically means that going a next step -- and I'm not criticizing the existing work, I think the existing work is great -- but one would need

156

information about not just patients, where they live and where they go, but who their employers are.

That would, I think, allow us to get some traction and make some progress to address the issues Greg has been raising. At present, I don't -- I don't think that has -- that has been done, but I -- all I'm saying is -- and while I'm not saying, okay, that is the research agenda, I think with -- with a bit of thinking -- I'm not saying this is trivial -- that one could identify, what would you need to do, what would you need to be able to do empirically, to be able to test a story that cross market mergers lead to competitive harm and distinguish that from one in which they don't?

And just to emphasize, this is really important. Obviously, we don't want socially, nor do we want the agencies, to go after mergers that are benign or beneficial, right? That's bad for everything. We want mergers that are beneficial to happen, and mergers that are benign, we certainly don't want to get in the way of any of that kind of thing, and the agencies don't either. So I think it is very important to try and get at that. And actually Matt's got some evidence that some mergers that go across markets can generate some real savings.

MR. BRAND: Matt or Matt, do you want to --

MR. SCHMITT: I guess just to speak to the concavity or convexity point, you know, my reading of the literature is that there's been a lot of focus on, you know, "must-have hospitals," that there are certain hospitals you just have to have in your network, and to the extent that there's a must-have hospital in Market A, a must-have hospital in Market B, and you really need both, that's convexity. That's -- I mean, they're complementary. If you don't have one, you don't have anything.

So I think, you know, actually generating concavity, I think it's definitely, you know, not clear that that's actually the structure of the payout.

MR. LEWIS: And it's not only the must-have hospitals, it's -- why would any two hospitals in far away markets be substitutes even for an employer with employees in both? So there's -- the linkage could be there, but it's not clear the direction of the linkage to me.

MR. BRAND: Let me throw out one further question. So if the -- so as described in Greg's work and in Dafny, Ho, and Lee, the basic notion of concavity here is payers negotiating with a set of

hospitals. If it loses -- one hospital's an employer with, say, employees in many areas, many of these areas would likely hang with that insurer, but if it loses two or more, then it's less likely to hang -- to stay with that insurer, so more likely to substitute away to another insurer.

One thought that's occurred to me is that it's quite -- it seems quite intuitive to turn that around. I think this kind of relates to what Matt said. It seems plausible to me that if you're talking to a health plan that is marketing its product to an employer with employees in two different -- two cities that are quite distant, and if that employer -- you know, the employer has to be -- has preferences over hospitals in each city, that the insurer may be thinking, you know, if I'm going up against three other insurers with both of these hospitals, if I don't have either one of these hospitals, I am extremely unlikely to win that business.

MR. VISTNES: And I think that kind of theory is -- that was really the heart of the theory that we tried to develop in our paper, and one of the things we found is that, again, it depends a lot on the assumptions, and the intuition here is, sort of going on what Keith is saying, the notion is a health plan

is saying, well, an employer is going to offer maybe just one or two different plans. If mine is not sort of the most attractive -- you can think of it certainly in the extreme -- if an employer is only going to be offering one plan, then that plan has to cover all the different islands in which that employer's employees live, and so if I get a hole on one of the islands, the employer can't offer it because it doesn't cover some of the employees.

Now, the more health plans the employer is offering, the more scope there is for me to have a hole in my network, because for any of the employees who don't like that health plan, because it has a hole on their island, they can pick another. So that sort of gives some wiggle room for the theory. But then you can also sort of think, is it going to give us convexity or concavity? Does the second hole hurt more or less than the first hole?

Then you can think of it in the following context, is let's say that all these health plans are kind of neck-in-neck, almost identical. In that case, my very first hole is going to put me at a competitive disadvantage relative to everyone, that first one knocks me out of the market. After that, you know, who cares? I'm already out of the market. The second

and third hole don't matter. I've got convexity as opposed to the concavity.

The flip side is, what if my health plan is fantastic? Everybody loves me. I can suffer this hole and they're still going to want me. I can suffer the second hole; they'll still want me. It's not until I get three or four holes that my superiority comes into question, and it's that fifth hole that really hurts me. Then I've got concavity.

So we've got -- we're back to, I'm going to keep calling it, the possibility theorem. How's that going to help me in a merger? How am I -- I guess in principle, but it's tough. It's tough to figure out when this is going to be a problem or not or, frankly, if there is the real theory driving it.

I won't say it now, but I think one of the other things we can talk about is, what are some of the other possible theories motivating some of this behavior? Because there are a couple of other sort of very different sort of potential explanations for what we're hearing. Maybe we're just on the wrong track.

MR. BRAND: Okay. I think we should probably move on to the bargaining weight. Maybe we will come back to other notions of convexity in the questions and answers.

161

So the next topic we'll turn to is the -- is whether the merger may cause a shift in the Nash bargaining weight, so a shift in how the joint surplus that's generated by an agreement between a provider and an insurer is divided between them. So my questions may include:

What are the likely interpretations of such a shift in terms of what determines the Nash bargaining weight in the first place? And how is the merger changing that? Is it just bargaining skills? Potentially something else? What is the likely effect on economic efficiency if that's what's going on? And, finally, could these -- if this is what's going on, could such be viewed as antitrust violations?

So I'll ask Matt Lewis to start us off, and then others can chime in.

MR. LEWIS: Yeah, okay. I mean, the important thing, given the discussion we've been having, the important thing to stress here is that the theory that we've sort of suggested in our papers as being potentially relevant is based on this bargaining -- bargaining weight is completely separate from these issues of concavity/convexity. It doesn't require any curvature in the profit function of the insurer. It's a totally -- you know, it's a totally different

162

mechanism, and it -- it also raises interesting questions because the potential conclusions from that kind of change are very different given that, you know, in the standard bargaining model, this weight just describes the split of the available surplus.

So do we care about how this surplus in the contract is being split between the hospitals and MCOs, if that's a transfer between those two, that maybe it doesn't have efficiency effects, but that's only for -- you know, for that particular contract, that may be true, and in the long run, there might be a lot of other effects as far as effects in the insurer market, which is why modeling the insurance market is important.

We may have, you know, a change in competition in the insurer market and an increase in pass-through to the premiums as a result of this. So I think these are all the interesting questions that come up here.

There's a separate issue which is kind of more of an empirical identification issue, which is that if you try to model these -- this bargaining power, it basically becomes the residual for anything that's not modeled in the bargaining position. And so you can get into a trap of finding a change in bargaining power when, in reality, you have just sort of left

163

something out of your -- out of your model of the bargaining position, and that's exactly why understanding better what drives the shape of the profit function of the insurer is so important, because if we really want to measure bargaining power, we need to -- we need to also -- we need to perfectly capture the bargaining position.

But having said that, I'll just quickly say that, I mean, I think without having a perfect measure for bargaining power, there's some, you know, evidence within what we have done which suggests that, you know, some of the existing theories on why you might see cross market linkages through bargaining position may not be as applicable to the situations where we do seem to notice some of these cross market merger effects, and that's why I still think bargaining power changes may be important here.

I'll let you comment.

MR. BRAND: Okay, we will go on to the next topic.

So on this next topic, I will also ask Matt Lewis to lead us off. So the next issue is that we may be, you know, bucketing up a wide variety of mergers into what we're calling cross market mergers, and it's possible that as we explore these mechanisms

164

that, you know, the mechanisms that are most important we will see only in a particular merger depending upon the specific circumstances of that merger.

So, again, I'll ask Matt to lead that discussion.

MR. LEWIS: Yeah. I guess what I would say here is I have in mind sort of an example where -- and this is an example of the kind of hospital we studied in -- or merger we studied in our papers. Think about a fairly large system, maybe 30 or 40 hospitals, acquiring a hospital -- a stand-alone hospital in a small town or small city somewhere. If you think -- you know, what do you think the effects of that merger might be?

If you argue that there's a potential for there -- for the merger to change the bargaining power, meaning change the bargaining sophistication of the hospitals involved, you know, do we think that that's going to happen to this acquired hospital? It seems likely that a stand-alone hospital might not have the same resources and experience in bargaining that a large system would, and maybe there's -- they can adopt some of those practices or use that information to better negotiate.

You know, do we think that that effect is also going to appear for the 30 other hospitals in the acquiring system? Do they get an increase in bargaining power? I don't see how that would be a significant -- a significant effect, and yet -- and now, if we're thinking about the different empirical papers, I think this is interesting, the Dafny, Ho, and Lee paper, they focus explicitly on measuring the cross market effect of mergers on the prices of these other hospitals, these 30 hospitals in the acquiring system.

So it's -- it's not surprising that we don't see cross market effects for those hospitals when we might see one for the acquiring hospital. I mean, I think there's a lot of -- that could also be true for some of these explanations of bargaining -- of bargaining position linkages, but I know it's important to compare the different sets of mergers that we've studied in these different -- in these different papers and try to understand, well, what does that reveal about where the sources of these price effects may be coming from?

And so I know I have sort of a strong opinion that I do think that in cases where small systems or stand-alone hospitals are acquired, this effect of potentially influencing bargaining power is -- is --

may be a big deal for those hospitals, but for some of these other mergers between large systems and other large systems, I don't -- I don't see why they would -- you know, they may or may not benefit as much, and we might be looking to other stories there to explain some of the findings there.

Again, Leemore or -- Leemore -- Leemore's paper also found that the effects were concentrated in her measurement on mergers that -- on hospitals that were located fairly closely but not in more distant markets, and I think that suggests something else, which we may or may not want to talk about, which is that maybe the patient market, as we're thinking about it, is not described or we're not thinking about it accurately. It may be broader than we might -- than we might otherwise think, so...

Did you want to comment on that?

MR. SCHMITT: I know we're not lawyers, but I'm curious whether this acquisition of a stand-alone hospital by a 40-hospital system, now we have better negotiators in place, is that something like -- you know, that the competition authorities should care about?

MR. LEWIS: Well, that's an important question, and it's one that I'd like to ask the competition

authorities, and I don't think I'm going to get an answer in public here, but -- yeah, I mean, it's a very important question.

Also, it's an important question because, whether that's true or not, acknowledging that those effects might be there may affect, you know, how we judge whether or not there are other types of effects. They certainly will affect the modeling if you have a structural model that does or does not allow those effects, that you may get those effects showing up in other places, and so I think it's important.

I don't -- I don't -- you know, I don't know -- I think we need a better model of insurance markets to know what this pass-through is going to look like, and even in that case, is there an efficiency effect that we care about and is there a reduction of competition? All the descriptions of the changes in the curvature of the insurance profit function, those very closely resemble the types of restrictions of competition that we look at when we look at in-market mergers, but this bargaining power thing is totally different and is not the same as a restriction of competition the way we normally think about it, so that's a very important --

MR. GAYNOR: Yeah. I mean, you can certainly get prices going up, and you can get harm to

consumers. You can get pass-through through the insurance market without there necessarily being harm to competition.

MR. LEWIS: It depends on the interpretation of it.

MR. GAYNOR: It depends. It depends. And I agree with you, I don't think at this juncture we know, right? We have these big effects -- which, again, that's valuable, we didn't know that stuff before -- but I don't think at this point we have a -- we have a good handle on that thing. And so, yeah, it's important for policy.

So if an effect like that occurs and if it's not through harm to competition, it might be of importance to policy, but it's not so obvious that it's an antitrust enforcement issue.

MR. BRAND: Final thoughts?

(No response.)

MR. BRAND: Okay. We're running a little late. We did want to touch on the point that Marty raised in his opening on potential broader implications of what we're learning in the literature. So let me throw that up, and, Marty, if you want to add to what you said or if any other panelists want to weigh in.

MR. GAYNOR: Just briefly, as I said, I mean,

169

this is potentially a very broad -- a broad issue, and so to the extent that what's -- you know, what's happening in healthcare markets provides an opportunity to try and really grapple with this and nail it down, then that's really useful, because it will give us an apparatus to start taking to other markets, not in a mindless way, of course, but at least to start thinking about that.

And, you know, one of the nice things about markets that are heavily regulated, like healthcare, energy, a few other things, is that there are a lot of data, because there are reporting requirements, so they can be good places to start trying to test some economic issues because of the richness of the data that are available, but I think we all agree that more thinking needs to be done at this juncture before we can figure out exactly or more precisely what's going on.

I think within healthcare, one -- we've talked about a few things that might be done. One avenue might be to pursue to look at -- look across product markets. Again, we have a lot of data on that. Folks are focused on geographic markets, and that's fine, but there's some other variation that's just sitting there in the data. Again, I think we need to think up

170

a way of sort of more precisely testing hypotheses before we just start crunching data, but I think there's some more stuff that can be done.

MR. VISTNES: I think I would maybe just add on to what Marty's saying. I think one of the real important things about the empirical work that's been done so far with respect to hospitals, what can be done even within healthcare, looking at, you know, the hospital physician or any of the other product market combinations, is if we're still trying to figure out what are the drivers of the theory, if we're still trying to figure out why is this effect occurring, then if we see the effect occurring, for example, across geographies, but we don't see it across different types of hospitals, or we see it between hospitals and physicians but not across different kinds of physicians, that will hopefully give us insights.

The -- looking at the data to find patterns, even if it is, in a sense, blindly looking at the data just to figure out what seems to be there, I think will help us figure out what is there or, alternatively, you know, decide that there isn't anything there, but more empirical work has got to be good.

171

MR. BRAND: Okay. With that, I think we should probably turn to questions from the audience.

MALE AUDIENCE MEMBER: Hi. So the possible sources of concavity, you know, the standard story is concavity is induced by competition between the hospitals, and now people have offered alternative sources of concavity, maybe, you know, the insurer would tolerate losing one provider, but losing two would be more than twice as bad, or maybe an employer for a similar reason would tolerate losing one provider but losing more -- losing two would be more than twice as bad, and those are sort of the alternatives to the standard competition story.

But another source of concavity that people don't talk about as much is just plain old risk aversion on the part of the managers of the insurance companies, right? So this is not concavity between the number of hospitals and profits. It's the concavity between profits and utility of the insurers, right? If you think that you are -- the profit loss if you lose one hospital is you would get a not great performance review, but the profit loss from losing two is you get fired, then that can be a real source of ordinary risk aversion that can introduce concavity, and that story seems at least as plausible

172

as the other ones. It does require that the insurer manage -- bargainer have a more concave payoff function than the hospital manager does, but that is entirely plausible. If you are -- if you are sort of the local manager of an insurer and you're facing the local -- the local sort of behemoth, it's perfectly possible that that's true.

So I would -- you know, I would encourage adding to the list of potential sources of concavity something that people are very comfortable with in other contexts, which is people are just plain old risk averse.

MR. BRAND: Anybody want to weigh in on that?

MR. GAYNOR: Well, again, I think -- I think getting data from insurance firms and the insurer market -- you know, this emphasizes that point of while it sounds kind of funny to think of insurance companies as risk averse? You know, there are very active reinsurance markets in which insurance companies buy insurance, so it actually does have some degree of plausibility on its face, but I think -- I think that could be certainly written down and specified, but then you're going to need the data to get at that.

MR. VISTNES: You know, I -- really quickly, to

173

sort of get to what are some of the other theories, a couple of the other sort of theories that have been bounced around is what I think of as kind of the crown jewel theory. Think of it in a typical department store mall, is you have got to have a couple of sometimes called like anchor stores. You have got to have a Cheesecake Factory or a William Sonoma or, you know, something else. None of those are substitutes, but they all sort of say, hey, I'm a quality provider. This is a good place.

Do I like the story, the theory behind it, but -- no, but can you understand how maybe it's going on in health plans? You know, I need a couple kind of crown jewels, and I can lose one crown jewel but not too many of them. Maybe that's a theory. You know, the other theory -- and I think, again, this is quite realistic -- is that people are not entirely rational. Unfortunately, again, we're seeing economists don't run the world and the world is suffering for it, but you have people who may believe, despite the fact that we're telling them you're irrational, your profits are linear, how many times do we need to tell you this, they say, yeah, but still, I kind of think I'm worse off losing two, and I think I'm a lot worse off.

If they believe that way, if they act that way,

174

it will generate all the empirical results we're seeing. It leaves us -- the enforcement folks and policy folks in the uncomfortable situation of, what do we do? It's -- it's kind of why we, in a sense, are uncomfortable with behavioral economics, because it doesn't make sense. How can we make policy based on stuff like that?

MALE AUDIENCE MEMBER: Yeah. I was wondering if I could hear a little more about alternative bargaining theories. I thought there was a little reference to why they work or why they don't work, and the reason is that Nash bargaining in this context is pretty new, it's pretty weird, and we do it mostly because it's feasible or it's a good place to start. As you say, the Nash bargaining parameter is a kind of residual, so you can look at it two ways, that if it changes, it means the true bargaining parameter changed, or it could mean that we're just sort of indicating that that's not the right bargaining model, right, that if that parameter doesn't stay fixed over time, that's just kind of a diagnosis.

And, you know, not thinking super formally, you can imagine very easily how people would think that coordinating bargaining across hospitals will let you do a better job. Now, Nash bargains are already

175

efficient and so forth, but if you think of, say, the multimarket contact collusion literature, where you can use a little bit of excess threatening power from one market to leverage, you know, a better collusive deal in Market B, nothing like that happens I think with Nash bargaining, but, you know, I think it's particular to the model, that really it's not a better -- this little hospital, when it says, oh, now I'm bargaining with this big hospital, of course, I'm going to get a better deal, right? I'm going to somehow get -- I am going to extract more somehow.

And, again, in Nash bargaining, the pie has already been completely and officially divided over there, so I don't -- I don't see how it works, but in the real world, I'm so sure things are so efficient and that there's not a little bit left someplace that can now be brought to bear on behalf of this new hospital.

MR. SCHMITT: Just to add something to that, to the extent insurer profits are meaningfully convex, then Nash-in-Nash bargaining can yield really strange predictions, which is, you know, just another problem on top of what you're raising.

MR. LEWIS: Yeah. I mean, I definitely think that it's worthwhile to try to figure out what these

176

other bargaining models would predict. I mean, the Nash-in-Nash model can be -- it can be generalized and it has been generalized to some extent in recent papers but actually still within this kind of general Nash framework, which isn't the best, I think, for this setting.

So on the other -- and also, the fact that you have this bargaining power parameter, which the theory gives you no insight as to how -- as to, you know, what determines this parameter, and we have a little bit of guidance from some of the -- some results in related bargaining models that may be information that plays a part in other things like this, but -- so we can use some of that intuition, but we know that this is a -- this is kind of an imperfect attempt.

You know, my position is just to say, you know, we don't know what will determine this thing, but it's -- it well could be heterogenous across hospitals, and maybe hospitals adopt that heterogenous bargaining power from their systems. Why does that happen? It could be information. It could be some kind of patient risk aversion -- you know, any of these results could apply, but I totally agree that a more -- a more realistic model of bargaining would be helpful here, but it's been a problem for us.

MR. GAYNOR: I'll just agree with everything that's been said, and, again, it's a possibility that the data and even what market participants had to say are telling us something, and we just need to go and think much harder about what the economic behavior is and what model that generates.

MR. BRAND: Okay. Other questions?

FEMALE AUDIENCE MEMBER: Just a quick question. It seems like what you're saying is, if I understand that correctly, you could have where one of these hospitals or both of them already have significant bargaining leverage or whatever because (off mic), so the merger is not necessarily going to change anything; it might or it might not.

I also heard what you were saying, Greg, at the very beginning, which is the theory and the evidence is in a state such that you can't really reliably predict when a given acquisition of a small stand-alone hospital by a large system is reliably going to result in some sort of anticompetitive effect.

So if we have this kind of uncertainty -- you know, one line of questioning is how do you go and work on and what kind of modeling to do in the context of a transaction, but let me ask to the panelists a

little bit broader question. There's a lot of move afoot in a number of states to think about regulating price, terms of access, other kinds of conditions on small, stand-alone hospitals, as they join systems, and as I think, Matt, you were saying, a very large proportion of hospitals are already in systems, although I think there's about 2000 that are still small stand-alone.

Let me just ask you on kind of that policy front, does that suggest we should be very, very careful about thinking of kind of across-the-board regulation or terms of access or pricing regulation on small hospital acquisition, because we might have some that could arguably lead to issues but others not?

MR. GAYNOR: Well, yeah, it's an interesting question, but I think, Meg, the -- it's not just a question about antitrust, right? It's a policy question, and there could be rationales for regulation that have nothing to do with antitrust, right, per se, but if there's a situation in which circumstances would change in a way that wasn't an antitrust problem but would really cause social harm, there can be a rationale for regulation. Very broadly speaking, of course, we should always think very carefully about any policy before undertaking it.

MR. BRAND: Okay. Any other questions?

(No response.)

MR. BRAND: Okay. Thank you very much. A very helpful discussion. Thank you.

(Applause.)

MR. ROSENBAUM: We will reconvene in about 15 minutes for the next paper session, so 2:00.

(A brief recess was taken.)

PAPER SESSION

MR. ROSENBAUM: If everyone would please be seated so we can get started. Thanks.

Okay, we are going to get started with the next paper session, which is chaired by Igal Hendel. The first paper is going to be presented by Paolo Ramezzana on contracting, exclusivity and the formation of supply networks with downstream competition.

MR. RAMEZZANA: Hi. So I will see how this works. So today I am going to talk about a fairly new way of looking at contracting in bilateral oligopoly with a particular emphasis on the endogenous formation of supply networks.

So before I start, let me give you the usual disclaimer, that whatever I say today does not represent the -- necessarily represent the opinions of the Federal Trade Commission.

Okay. So a lot of markets look approximately like this. You have some downstream firms -- I've drawn two here, R1 and R2, where R stands for retailers. These downstream firms procure differentiated inputs or products from suppliers, and the dashed lines you see there are potential supply contracts. And when these downstream firms has secured some supply contracts, they compete for consumers in the downstream market, right? So that's a typical bilateral oligopoly setting.

Now, some markets look like this. There is -- in some markets, all supply links are active or at least most downstream firms carry most products. Okay, examples of this are big box stores, like Best Buy, Target; online retailers, Amazon carries pretty much everything; and online travel agents that carry pretty much all flights, all airlines, and all hotels, right?

But other markets look different. So in other markets, some links are not active. So, in particular, the downstream firms may decide to carry different types of products. So a good example of this is the cell phone industry a few years ago -- it's sort of still the case, but especially a few years ago -- when the iPhone was launched, it was launched exclusively by AT&T, and that's -- you know, for four years, and that's the best known example, but it's not the only one.

Around the same time, the Google phone -- you may remember the HTC G1 phone -- was launched exclusively by T-Mobile, and also some LG models were launched exclusively by Verizon. So pretty much every

wireless carrier was offering exclusively some different type of handsets. And here you should think of the handsets, S1 and S2, and the wireless carriers as the distributor, R1 and R2, okay?

There are other examples, sport events in pay TV. Typically sport events are broadcast exclusively by a channel or a different channel, MVPD, a different platform, and lately, health insurance companies have started offering restricted networks, okay? So all the examples I gave you there involve some type of contractual exclusivity.

However, there are also examples, like the automobile distribution in the United States, where there are no exclusive contracts, because those are actually prohibited by law in the United States by a crazy maze of state laws that prohibits that, yet different car dealers typically specialize in different brands, okay?

So these are interesting patterns. So what are the research questions? What are the -- is there any interesting research question from a theoretical point of view?

So the first one is, what types of supply networks maximize industry profits -- that is, produce a surplus -- and what type of networks maximize

consumer welfare? The other question, which arguably is even more interesting, is what type of supply networks can arise as an equilibrium when firms engage in decentralized contracting, right?

So there is an old literature sort of addressing the first paper. Some of you may be familiar with a paper by Bazan and Perry (phonetic) a long time ago talking about that, but we really don't have a full-fledged model addressing the second question, what are the equilibrium networks? So here I develop, I present a model of bilateral contracting in which firms can use transfers to induce other firms to enter into a relationship with them, okay?

So this model combines two streams of literature. One is the literature on the formation of economic and social networks with transfers, so Bloch and Jackson is an example, there's a lot of work by Jackson and others on this; and with the literature on vertical contracting, and there's a few famous papers there, okay?

So this framework allows me to identify, to study a few factors that may actually influence the -- affect the structure of supply equilibrium networks. So the spectrum includes the degree of supplier and retailer differentiation; the mode of downstream

competition; Cournot/Bertrand, how intense it is; the availability of exclusive contracts; and the firm's ability -- or actually, in the context of this paper -- inability to commit to the terms of the contracts. Okay? So that's the broad picture of what I do.

Now, one may say, well, but we do have a framework, which is actually very popular in IO at the moment, which sort of looks at contracting between multiple suppliers and multiple retailers, and that's the Nash-in-Nash bargaining framework, right? And there you can see some of the papers in this literature.

I particularly want to draw your attention to a recent theory. There is a paper by Collard-Wexler, et al., that provides some theoretical foundations for Nash-in-Nash, and more to the point of this, it provides a very nice discussion of the assumptions on which it is based and on the limitations of those assumptions.

Okay. So what are these assumptions or what is this approach? So the first thing to say is that Nash-in-Nash focuses more on the division of surplus between suppliers and retailers rather than focusing on the structure of vertical contracts or focusing on the structure of the networks that emerge. So that's a different type of question it's answering, right?

The other assumption is based on the contract equilibrium approach; that is, when two firms negotiate a contract, they take all other contracts as given, including contracts to which they themselves are a party. So if a retailer has negotiated a contract with a supplier, that retailer is not allowed in that approach to maybe modify its contract with the other supplier, right?

Now, it turns out that's not a big deal if all they want to do is predict the division of surplus, and, in fact, Collard-Wexler, et al., show that there is fairly general conditions. Nash-in-Nash bargaining gives you the same result, is a more general, multilateral bargaining -- strategic bargaining aid, okay?

It is, however -- it is, however, a problem for what I want to do here. To see that, look -- follow the following example. Consider a supply network, a contract equilibrium supply network, in which everybody trades with everybody, okay? And now consider a deviation in which a retailer, R1, approaches S1 and asks for exclusivity. It asks S1 not to trade with R2, right?

So in the contract equilibrium approach, that's all he can do. He cannot go to S2 and try to modify the other contract. He can only modify one contract at a time, because he has to take it as given that S2 continues to trade with R2 in that approach, right?

Assume that this deviation is not profitable. The parameters are such it's not profitable, right? Well, there is another deviation if one looks at the marginal approach, which would be my approach, in which R2 could approach both suppliers at the same time and ask both suppliers to be excluded within, you know, excluding R2. That deviation might well be profitable even if the one in the middle is not. So by focusing on Nash-in-Nash -- on contract equilibrium, using Nash-in-Nash bargaining, you may be missing something, okay?

So another assumption that Nash-in-Nash uses is that given an exogenously given set of links or networks, every bilateral negotiation, every link, it's assumed to yield gains from trade. An implication of that is that the only equilibrium -- the only possible equilibrium is all links active.

There are some recent papers by Ho and Lee that discuss these issues. I'll talk about them at the right point in the presentation, not now.

Finally, Nash-in-Nash typically simplifies the structure of vertical contracting. It either assumes that payments between suppliers and retailers are only lump sum, without any margin on input price, or goes to the opposite extreme, that they are linear, okay?

Okay. So the approach I propose today, before I give you the model, improves on this along the following dimensions: It allows firms to optimize across all their bilateral relations at the same time, to modify all the contracts at the same time. It allows firms to use nonlinear contracts with a fixed fee and a marginal input price. It allows firms to enter into and actually compete for exclusives, okay? And, especially, it's sort of able to generate -- to give predictions on the endogenous emergence of supply networks or a type of supply networks, right?

So these are the advantages -- oh, sorry -- but my approach, to be fair, also has some drawbacks. So Nash-in-Nash gives point predictions regarding the division of surplus. It will tell you exactly -- you know, it will give you a price point. My approach, as you will see, would only give you ranges for the transfer. It would only give you the bargaining set of the transfers. One can get quite a bit of mileage out of that, as I will show you, but to be clear,

that's a drawback. For applied work, one needs to do a bit more, okay?

Okay. So let me give you a sketch of the model. There are more than two suppliers, in this by S, more than two retailers, in this by R, and, of course, if all of these firms are active, you have S times R differentiated products.

The model evolves in two stages. In the first stage, all firms engage in simultaneous contracting without public commitment. So it's secret contracting. Firms cannot commit to the terms of the contracts, okay? And once all contracting is done, in stage two, retailers with at least one contract engage in downstream competition. It could be Cournot, Bertrand. I actually address both, okay?

Now, stage two is completely standard here, okay? So let me talk about stage one a bit. So in stage one, each firm I submits a contract proposal to each firm J on the other side of the market. This is basically an extension of Bloch and Jackson, 2007, to vertical contracting.

Each firm I submits a proposal to -- a contract proposal to each firm J on the other side of the market, so all firms submit simultaneous proposals to other firms, and the contract proposals contains three elements. One is a lump sum or a fixed fee, if you will, to be paid by the retailer to the supplier; the other is a unit wholesale price; and the third is a set of exclusive clauses, if any. There could be none, okay?

Now, if the proposals that two firms, say supplier S and retailer R, submit to each other are consistent, then these two firms enter into a contract, and the supply link is formed, okay? And a proposal that is consistent, if both firms name exactly the same wholesale price, exactly the same set of exclusive clauses, and the retailer offers a lump sum which is at least as large as the one that has been demanded by the supplier, okay?

Now, a model like this is replete with coordination failure, vertical coordination failure, horizontal coordination failure, so I'm not going to even go into that. So there's a ton and a half of Nash equilibria with different networks. So Nash equilibria is really not the right -- I mean, this is on purpose. I did the model like that on purpose. Nash equilibria is really not the right concept here.

So instead I rely on coalition-proof Nash equilibrium. So the nice thing of coalition-proof Nash equilibrium is that it allows players to engage

in prepaid, nonbinding -- and the nonbinding part is important -- communication, right? So to be clear, it can't be used to enforce collusion, because firms, you know, can commit to what they discuss, and so every type of agreement that is reached must be (indiscernible) compatible. So you still have a lot of space for competition and division and all of that, okay? It's just a way to eliminate silly coordination failures.

And going a bit more into details, the outcome is a coalition-proof Nash equilibrium if there is no deviation by any coalition that leaves all the members of the coalition better off. And that's not the end of the story, though, because this deviation must, in turn, be robust to follow the deviations, okay? There must be in other deviation from the division, so on and so forth.

It's very similar -- to keep it simple, it's very similar to subgame perfection, okay? You can find a profitable deviation, but once you get there, it -- you know, you may want to do what you set out to do, okay? So it's just some consistency.

Okay. So how can one use this to solve the model? The model can be solved in two steps. First, for any network g, you must find a profile of

wholesale prices such that there is no dev -- such that in the same network, they cannot re-arrange the wholesale prices and make profits, okay?

Now, without public commitment, with secret contracting, there is obviously opportunism. So the only wholesale price with that characteristic is wholesale prices equal to marginal cost, okay?

Of course, if firms were able to commit publicly, then the wholesale price would be greater than marginal cost, and actually I am working on a related paper, but -- or if you must give firms incentives to engage in an ongoing investment, if there is a problem or a hazard, again, the wholesale price is above marginal cost, right? But in this stylized model, the wholesale price is equal to marginal cost.

Now, this is not new. It's a very standard, well-known result from the vertical contracting literature. All I do here is to extend these inside to a much more complex environment, with multiple suppliers and multiple retailers, and to a different equilibrium concept, coalition-proof Nash equilibrium, which, by the way -- you can read the paper on that -- but it turns out to be very convenient, because it solves existing problems that have been identified by

Rey and Verge in 2004.

But all I want to say here is that the result on wholesale price is a key ingredient to what I do, but it's really not a big contribution on the paper. It is just an ingredient, okay? So the main contribution of the paper is characterized in equilibrium networks, right? So in this model, when network g is in equilibrium, if there exists at least one profile of transfers, tg, such that there exists no deviation from the network g, then it's profitable for all the firms and it is self-enforcing, okay?

Now, to be clear, it's enough for there to be one transfer for g to be in equilibrium, but typically, there are many possible transfers that support an equilibrium g, and so what I'm doing here, I'm just really only characterizing the bargaining set, the set of terms, okay?

Now, let me go back a second. So in the paper I discuss some general methods for verifying whether a division is profitable and self-enforcing. I don't have time to go into details here. Let me just show you very quickly just what the intuition is.

For example, a deviation from a network g to a network h is profitable for a coalition Z if the gain in gross profits that we generate, the change in the

193

gross profits produced because of those effects I was talking about before, are greater than one, are greater than the change in the transfers received by suppliers.

So if the deviation involves dropping some retailer from the network, the suppliers were getting transfers from those retailers, right? And so they are going to lose that if they drop them. So that needs to be taken into account, okay?

And analogously, you have to take into account the fact that if the deviation were to drop in some suppliers, then the retailer no longer pays to those suppliers, okay? So that's just intuition. You don't need -- by the way, you don't need to remember any of this for the rest of the presentation. It will become very intuitive in 30 seconds, okay?

So, but that's one result, and all the complication in the paper, I'm not going to go through this now, but it's -- if you find out that there's no profitable deviation from g, then you're done. Answer. g is an equilibrium; in fact, it's a strong equilibrium, right?

But if you find some profitable deviations, then you still need to check that those are self-enforcing, right, from the logic I said before,

194

and that's complicated, because when you check for change of deviations with transfers like here, you have to take into account that a transfer that can make a deviation profitable depends on what the transfers were in the previous allocation, and so there are chains, and that's a very -- one of the very complicated aspects of solving this, okay? So -- but this is a general approach. As I said, you don't need to remember much, just to give you a flavor.

Now, let's move to the example that I will use throughout the rest of the paper. So I look at a bilateral duopoly with two suppliers and two retailers, and linear demand, okay? In a bilateral duopoly, you can have a number of networks. The ones you see on the screen now are networks that actually will arise in equilibrium, right?

And the networks you see here are all the networks that want that equilibrium, but they feel it's impossible. And actually, out-of-play networks matter here to find equilibrium. So they are on the network that can arise.

Now, demand. I used linear demand, and it's a convenient thing because it allows me to parameterize product substitutability, using a prompt a, and retailer substitutability, indexed by b, separately,

195

okay? And the nice aspect of this is that you can represent all the results in a unit square, with a and b on the two axes, okay?

So one can use this model to answer the first question. What supply networks maximize industry profits? So this Bertrand competition, let's start from the top, okay, where b is close to one and the four retailers are close substitutes. In that case, industry profits are maximized by downstream monopoly, and it's very intuitive, because if b is close to one, retailers are very close substitutes, there is a lot of competition to kill off, right, and it's not very costly to kick one retailer out because they're very similar, okay?

Where retailers become more differentiated in the intermediate space, right, it starts being costly to eliminate the retailer. So you want both retailers to be active, but competition is still pretty strong, so you want them to carry different products. And eventually, when retailers become very differentiated -- there isn't much competition to kill in the first place, and it's very costly to keep a link out, so you want to have all links active, okay? So that's the network that maximize total industry profits. That's for Bertrand, Cournot, it's very

196

similar results, okay?

Now, let's move to the real question. What are the exclusive contracts that emerge -- what are the networks that emerge in equilibrium when we don't have -- let's start with a case in which we don't have exclusive contracts arrangement. So if you don't have exclusive contracts and you have Cournot competition, then the only possible equilibrium is with all links active. Everybody trades with everybody. It's very intuitive, right?

If I'm a firm, I can't prevent my counterparty from dealing with somebody. If that counterparty finds it profitable, they will do it. And so everybody trades with everybody. It's very intuitive, right? And so this makes things agreeable, and profits are not maximized.

However, it's not general. If the -- you have Bertrand competition -- that should be grayed out, the top edge should be gray, it doesn't show up well. When retailers are close substitutes, pairways exclusivity is the equilibrium, and the intuition here is the following. In providing exclusivity, the two retailers carry two different products.

Now, imagine retailer one is thinking of also getting product two. Now, the good thing is that it

197

gets variety, but the bad thing is that if he gets good two, it gets the same good that its rival is carrying, so it becomes more similar to its rival. That will instigate a reaction from the rival. That will instigate the rival to lower prices and start a price war.

So it's completely possible that retailer one just decides to forbear. He could get the other product, but he just doesn't get it, okay? And that's how you can have provider exclusivity even in the absence of an exclusive contract, okay?

Okay. Now, let's see -- I have to go fast, obviously, given time constraints. Let's go to the case with exclusive contracts, all right? So you have to adjust the framework a bit. It's actually very tricky, but I am not going to bore anybody with that. So I won't tell you what you need to do to the framework, but once you make the framework consistent and ready to go, these are the results.

Let's look at Bertrand competition. These are the three areas in which different networks maximize total industry profits. There's not equilibrium. That's what maximize total industry profits, okay? So downstream monopoly provides exclusivity, and all links active from top to bottom.

198

So what happens? Now, first of all, when downstream monopoly is in equilibrium -- sorry, when the downstream monopoly maximizes industry profits, it can always be supported as in equilibrium. Going to the opposite extreme, when Bertrand comp -- when retailers are very differentiated in the lower part of the graph, all links active can be supported as an equilibrium in most part of the -- no, in the grayed part of that region, and for intermediate levels of retailer differentiation, the middle area, it turns out that provider exclusivity can be supported as an equilibrium when the suppliers are very differentiated, when a (indiscernible) differentiated, which makes a lot of intuitive sense, because the reason why provider exclusivity increases profits is that it allows retailers to inherit the differentiation of suppliers, right? And so you would expect it's more likely to be in equilibrium when suppliers are very differentiated, okay?

So what I've shown you here are two-strategy equilibria. The bad news is that in the white area, it really doesn't exist, the two-strategy equilibria. There could be mixed-strategy equilibria, and that's not even sure with this type of equilibria, but I don't want to -- okay? So that's what happens with

199

Bertrand. Cournot is very similar. You know, it's different regions, but it is the same, okay?

Now, before I conclude in the next three, four minutes, let me talk about some implications of these analyses, okay? So the first one is very straightforward given this model, and it is that exclusive contracts in this model always reduce welfare. That is -- well, let me say it better.

When exclusive contracts are actually adopted in equilibrium, so in that area where they actually are adopted, and when they actually cause the equilibrium to switch, which is not everywhere, but when they have an effect, they always cause the equilibrium to switch in the direction of less variety, right, and of less competition, of higher prices, and that means it's bad for welfare. We all know there are all the potentially positive effects of exclusive contracts, but in this model, they're bad, okay?

Now, much less straightforward and much more interesting, in my view, is the effect of exclusive contracts on the distribution of profits between suppliers and retailers. As I told you, here I can only predict ranges, right? And the upper and lower bound of those ranges are determined by the credible

200

deviation, the credible threats available to suppliers and retailers.

In particular, let's just focus on t-upper-bar, just for the sake of example. t-upper-bar is determined by the credible deviations available to retailers, and the idea is because if suppliers want to raise -- they want to get a transfer which is very high, if a supplier wants to do that, eventually, a retailer will just kick him out, right? The retailers would kick him out if he is too high, right?

The ability of the retailer to credibly kick him out determines how high the t can be, and the term is t-upper-bar, and (indiscernible) for t-over-bar.

Now, it turns out that in this model, the availability of exclus -- notice this idea of the availability of exclusive contracts. They don't need to be adopted in equilibrium. The sheer fact that they are available changes the credibility of deviations, and it makes it more credible for suppliers and retailers to exclude somebody on the other side of the market. So it affects the outside options.

It turns out that in this model, it affects the outside options of the retail -- of the suppliers much more than those of the retailers, and I can discuss that later during questions, but the availability of suppliers can kick a retailer out and implement a downstream monopoly that's very profitable relative to the alternative, what the -- relative to what -- the recourse retailers have.

And so when -- in this model, when exclusive contracts become available, they make suppliers unambiguously better off and retailers unambiguously worse off. I would like you to note here that this approach is very similar to -- so this is very similar to something that Bernheim and Whinston found in a 1998 JPE paper on exclusive dealing. It's different from the approach taken in Ho and Lee in a recent paper and in the paper you may have seen by Eli Liebman. In that case, in those last two papers -- so, first of all, they don't have downstream competition, so -- well, Liebman has it but doesn't make much with it, and Ho and Lee assume there is no downstream competition, so they focused on a different issue.

But basically in those papers, the idea is that retailers -- that's health insurance companies in their model -- can commit to exclude, ex post, one or more suppliers, one or more hospitals, and by committing ex post -- by creating artificial scarcity,

they will induce hospitals to be more aggressive and get -- so they will get better terms, but in order to obtain that, they actually do need to exclude somebody on the equilibrium path and cause some damage. That's not what happens in my model.

Finally -- and I'm almost done, just basically one minute -- but I want to talk about two papers, one by Lee and Fong, Robin Lee and Fong, and the other one by Rey and Verge, which ask a similar question. They look at what type of supply networks arise, but their approach is quite different from mine. They assume that firms first form all the supply links, the network, without being able to use any transfers at the network formation stage, and they can't even use long-term contracts. So they can just -- you know, they can just form the networks without compensating each other.

Once the network has been formed, then they Nash bargain, and Nash bargaining takes place under conditions of hold up here, right? What does that do to the equilibrium of the model? Well, hold up makes it more difficult for two firms that want to create a link -- and I assume that's jointly profitable -- to move money around to make sure that happens, because one of those firms is afraid maybe to be held up and

won't play along.

You can always delete a link unilaterally, so it doesn't do anything to the ability to delete a link. It only makes it more difficult to create links, hold up. So it tends to produce networks that are narrower. It's more difficult to withstand the network. And so the two figures you have there on the left is my approach with transfers, and as you see, the bottom part is all links active; the upper part, there is some exclusivity. The right side is one that occurs in Rey and Verge, there's more exclusivity, right?

Nice approach, interesting in some markets. I don't think, though, that it's very realistic in markets with large firms, like the deal between AT&T and the iPhone -- sorry, AT&T and Apple on the iPhone, there were big payments probably up front, right? And it's also not very suitable to study in exclusive contracts, because no firm would commit to exclusivity if it can't be compensated, right?

So, in conclusion, I developed a new way to look at bilateral contracting in -- bilateral contracting in bilateral oligopolies. These identify some potentially important factors to determine the structure of supply networks, but so far, it has

focused more on the division of surplus than -- more on the structure of networks and contracts than on the division of surplus.

So possible next steps would be to do more work on the division of surplus. That's really complicated, but that could be a step. And the other one is to find a way to empirically implement this, to simplify it and empirically implement.

And the other thing one could do is study markets where firms can publicly commit to the wholesale prices, and I'm doing that in ongoing work.

That's it. Thank you.

(Applause.)

MR. ROSENBAUM: The discussant is Ali Yurukoglu.

MR. YURUKOGLU: Okay, thank you for inviting me and thank you to the organizer --

MALE AUDIENCE MEMBER: If you could get closer to the microphone.

MR. YURUKOGLU: There's a lot of -- it was very interesting to read this paper, a lot of rich economics. Let's jump right in. I was going to start by motivating with some examples. I think Paolo did a good job of that. Let me mention one or two more that he didn't mention.

205

So you see this with -- these exclusive deals with department stores and clothing brands. For example, Target will do these collaborations with high-end designers that are exclusive to Target; also, soft drinks in restaurant changes. In many cases of bilateral oligopoly, we see interesting cases of incomplete supply networks.

This paper is really about defining equilibrium notions that will get you those interesting cases and trying to generate networks like this in markets where buyers and sellers have market power, payoffs are interconnected across negotiations, and contracts are potentially complex, not just about price. And like he mentioned, it's really sort of combining two different theory literatures, one on vertical contracting and one on coalition-proof Nash equilibrium.

Okay, so I am going to sort of have a high-level comment about both of those, which I'll go into the details, but -- so a lot of what makes this go is the assumption of secret contracts, okay, and flexible contract spaces. That's what gets you -- it makes it easy to solve the pricing equilibrium, okay? So you get wholesale prices which are equal to the marginal cost of production.

206

And so my comment about that is going to be, well, how do you deal with the fact in reality that we see very often linear prices above wholesale cost, and if you build that in in a natural way, would that change the results?

And then I have some comments on the -- pointing out some trade-offs between thinking about using coalition-proof Nash equilibrium or something like a Nash-in-Nash equilibrium.

So I'm going to refer to Nash-in-Nash as Horn and Wolinsky, they're the same thing, basically comes out of this Horn -- this paper by Horn and Wolinsky. There's a common misperception, I'd say, that Horn and Wolinsky has nothing to say about equilibrium supply networks. It does. I'll show you a simple example now, which is basically some supply networks can't be part of any Horn and Wolinsky equilibrium.

So if you just want a very simple example that generates this, imagine you have two upstream -- identical upstream manufacturers and a single downstream retailer, okay? The network where only one manufacturer provides to that retailer can't be part of any Horn and Wolinsky equilibrium, okay? That's because you have to think of the uncontracting party.

There's a negotiation problem between those two

207

that's not being solved, okay? So if they have identical costs -- the manufacturers have identical costs and you propose an equilibrium with only one link, okay, if that price in that link is above cost, there's an incentive to sign a contract with the other manufacturer. So that can't be in equilibrium, okay?

And if that price is at exactly equal to cost, that pair actually has an incentive to deviate the price upwards when the other firm is not there, okay? So Horn and Wolinsky does, in fact, give you predictions about equilibrium supply networks, okay? So that's sort of a starting point.

Now, Horn and Wolinsky has its own warts, okay? So I have heard Steve use the adjective "weird," also I've heard "schizophrenic," or "unnatural," okay, lots of colorful language. It's true, Horn and Wolinsky only looks basically at pairwise deviations, and some of those are you might think extremely unrealistic, because they're holding your own company's contracts fixed when you're thinking about what would happen if we were to sign a different contract with another party, okay, and that feels a little unnatural, though I -- I'm not going to get into it here, but there are some very good theorists who think of that as a feature, not a bug, and perhaps on that -- perhaps

208

less unrealistically, it doesn't deal with deviations that involve multiple firms, okay? So that's what Paolo is getting at here with the coalition-proof Nash equilibrium.

Okay. So the main difference here is that instead of Horn and Wolinsky only looks at these pairwise deviations, the coalition-proof is going to look at multilateral deviations. Let's just think about the trade-offs in using that for analyzing real markets, okay?

One thing is that when you only have two sides of the market and you're thinking about a multilateral deviation, that's necessarily going to involve two firms on the same side of the market, okay, which is going to lead to issues of horizontal coordination, like do we think that these firms actually can make those deals?

I would like to see an equilibrium notion that, if it wants to get at multilateral deviations, it sets it up in a way that, like, the communications only go through the vertical channel, that it doesn't involve firms, you know, who compete with each other coordinating their deals, you know, you take -- you take that input, I take this input, and we'll agree not to go on each other's territory, because they might for legal reasons not want to do that.

The other thing is -- so these words like "schizophrenic" and "unnatural," "weird," I think can be applied very well to the coalition-proof Nash equilibrium as well, okay? So it's got a wart, which is that the deviations you have to check for only have to be immune to further deviations within that pair, okay, where you might think, well, if there's a profitable deviation by a set of three firms, okay, it might be that once that deviation is made, there is now a deviation in that world consisting of some sets of those firms and a third party who wasn't part of the original deviation, okay? That's not ruled out in coalition-proof Nash equilibrium, okay? So that seems a bit schizophrenic to me as well.

Okay, and another complaint about Nash-in-Nash is, well, what's the game, the noncooperative game that gets you there, okay? Is this sort of a similar question here?

Now, the benefit of Nash-in-Nash, which I've mentioned, is tractability, and that's, I think, a clear benefit here, which is for the analysis we restricted to two-by-two for computational reasons, right, the number of combinations you have to check gets large, so I would be sort of curious to know how

well this performs when you have bigger networks, like realistic networks in terms of size.

I think estimation, you could probably -- it might be one of those cases where, like, it's -- you can estimate it, but it's much harder to simulate it, okay, because you could just use the necessary conditions for estimation, but this is all just to say this is worth looking at. It's not -- like, this doesn't in one fell swoop get rid of all the problems of Nash equilibrium. It's not like a Pareto improvement, but it's something to add to our toolkit. It might be more applicable in some industries than others.

Okay, along similar lines, I have seen a bunch of papers recently that sort of take standard supply/demand models in IO, that's BLP Demand, Nash pricing at the downstream level, and they try to generate interesting supply networks by playing around with the rules of the contracting game. When I think there's actually -- like, there's an alternative, which is you could try and play with the supply and demand models to generate different incentives that will lead to different supply networks, okay?

So, like, most of these models have linear cost functions, okay, whereas you think certain types of

nonlinear cost functions will lead to exclusive dealing, okay, so that if by shipping you -- like think in the hospital case, if I ship you a lot of quantity by putting you in a narrow network, then the hospital knows it's going to get a lot of quantity, okay, and if the hospital's cost curve is concave, you're moving the hospital to a flatter part of their cost curve, lower marginal costs, so you should expect better prices in that case, okay?

You know, costly capacity for the retailer might be a reason you don't stock every item. Nonlinear pricing by the downstream firm, in a lot of those narrow networks, the insurance company actually has a deal with the hospital that's not in the narrow network, and they use that hospital in other products, okay? So they are negotiating. They just don't offer it to -- in certain products. That seems more about product design at the downstream level than about, you know, some weird trick on the contracting game.

One-stop shopping by consumers, like why does Target have those exclusive collaborations with designers? Okay, you know, there's models out there that say if it's hard to observe prices, but you can observe what's being stocked, like they do a promotion saying we have this collaboration, and you have

one-stop shopping, okay, then that's a way -- that type of exclusivity is going to be gen -- is going to be generated without any sort of playing around with the fine details of contracting.

Okay. So I think it would be useful -- if we want are models of incomplete supply networks, I think there's a lot of room still to play with demand and supply conditions rather than details of the contracting game.

Just as a last comment, so I mentioned this at the beginning about -- so a lot of the analysis is simplified I think in a very pragmatic way by assuming that contracts are secret and there's a flexible contract search, so a two-part tariff is enough, okay? In those models, in any equilibrium, the price that the manufacturer charges the retailer is going to be equal to the manufacturer's cost of production, okay?

This is very robust, goes back to Hart and Tirole. It seems natural because there's nothing really preventing firms from using flexible contracts. The problem is, in reality, we see linear pricing above wholesale costs sort of all the time, okay, cable TV, music streaming, certain medical procedures, something like basic inputs for basic industries. So I think these models are missing something that leads to more linear contracting, because there's other stuff going on that is pushing these firms away from that benchmark of getting linear prices at wholesale costs. And I'd be curious to see, when you put those in, sort of how much you can still do and whether it would change the results or not.

Okay. So to wrap up, it's a really interesting paper wrestling with really important issues in antitrust and IO. I think, you know, very theorists look at this is very fruitful right now. It combines insights from contracting vertical relations with coalition formation theory. You know, a nice part of the paper is it predicts a wide array of supply networks, which I think is great.

I'd like to see a little bit more about what this coalition-proof can do that Nash-in-Nash cannot and whether it's worth the computational -- you know, Nash-in-Nash with -- allowing the analyst to play with the supply and demand model and whether, you know, those benefits are worth the computational costs or not.

Thank you.

MR. ROSENBAUM: We have time for some questions.

MR. RAMEZZANA: First of all, maybe just really

quickly, Ali, really an excellent discussion. The only thing I wanted to say is that, absolutely, wholesale prices are different from marginal costs generally. I'm actually working on a paper now in which there's public commitment, and so they will be greater, but, yeah, there are a lot of circumstances in which you have double moral hazard, and that's the case, so that's a great comment, and I agree with that.

MALE AUDIENCE MEMBER: I believe this is taking a slightly different angle, but I've found that one of the issues with Nash-in-Nash is that when you're thinking about the disagreement payment, you're not allowing to have a next round in which you renegotiate. Even if a contract fails with one party, then you can renegotiate with another potential party. So like if a hospital does not agree to something and maybe you are going to divert and change your price, I mean, that will change the bargaining with the -- you know, with other hospitals.

I heard there are sort of recursive concepts of Nash-in-Nash. I mean, are you familiar with that? I mean, is this --

MR. RAMEZZANA: Well, so what he's saying is that they could recontract the outside options in some

sense. If we did agree, I can go back and re-optimize my outside option, not only -- no, I should look into that. Yeah, no, I mean, to be honest, I -- no, I'm not familiar with the -- with those more in-depth treatments. I should look at it.

There are -- there are a number of other equilibrium concepts one should explore, and eventually I may get to that. Also, sort of in the literature on coalition-proof Nash equilibrium, coalitional equilibria, there are -- equilibrium is -- there's a book by the by Bloch and Dutta, you know, if you look forward, because here basically you -- if coalition deviates to a certain outcome, then it's completely nearsighted. It's not looking at the fact that once they get there, they maybe deviate farther. They just do things one step at a time, and they just get there and say, okay, this doesn't work, something else happens.

The smart people would usually think, okay, if we go there, this is going to happen, and they are going to look at the endpoint of this. So that's not what I did here, and that's not how CPNE works. So I'm not sure it addresses your question, but generally, I think there are -- I agree with you, given -- this has been a lot of time already, but

given a bit more time, one can try to figure out more refined equilibrium concept, which maybe make -- either make this consistent with Nash-in-Nash to some extent or anyway can help. Yeah.

MR. ROSENBAUM: Any other questions?

(No response.)

MR. ROSENBAUM: We will move on to the next paper. Thanks, Paolo.

(Applause.)

MR. ROSENBAUM: All right. Next we have Michael Geruso, presenting Contract Design: Evidence from the ACA Health Insurance Exchanges.

MR. GERUSO: First, thanks very much for having me. I'm really happy to be here at my first FTC Micro Conference. This paper is joint with Tim Layton and Daniel Prinz. Daniel's a graduate student, not on the market yet.

Right, so just to give a bit of background and orient you to what we're going to do in this paper, sort of at the heart of this is this fundamental tension in health insurance markets between offering consumer choice, trying to enforce nondiscrimination -- and we can talk about why nondiscrimination might be a good thing from a social planner's perspective -- and then, as a result of

217

that, dealing with selection.

So in the exchanges, which is what our paper is about, that's our market setting, but this is also true in Medicare, in privatized Medicare and privatized Medicaid, the rules of the game are that you need to enroll anyone who wants to join a plan, you can't charge different people different prices, and, in particular, you can't -- you know, so there's -- you can charge people who are different ages different prices, but you can't link premiums that people pay in these markets to their health status. That's something that's very popular among consumers.

So if you think of the recent debate over repealing and replacing the Affordable Care Act, the idea of preexisting conditions and coverage for those has come up over and over again. And so these regulations enforce a fairly intuitive sense of fairness in these markets, but they also connect -- you know, you can backstop all of this with very clear economic theory about insuring consumers against exposure to long-run risk.

The trouble with these kinds of regulations is that they also open the door for inefficiencies related to selection, and the reason is is that price

218

is just one of many potential screens in these markets. So you can say that you can't charge different people different prices, but what's much harder to observe and what this paper's going to be about is do you -- do you distort other aspects of the contract to try to keep certain people out of your plan.

Now, it turns out that there is a very widely implemented and standard solution to this problem, which is risk adjustment, and since probably only a minority of the people here are into health insurance, the basic idea behind risk adjustment is you want to give the insurer a payment that compensates them for the expected cost of the enrollee that they're taking on. So if you are going to take on someone with diabetes, then the regulator is going to take a payment away, is going to tax a payment away from plan that enrolls a healthy 25-year-old, and it's going to give that money to a plan that enrolls, you know, a 64-year-old diabetic.

And when that's working properly, at least under conventional wisdom, you are just exactly compensating expected costs, and all enrollees look equally profitable even though they're differentially costly. That's the basic idea. That's widely used in

219

Medicare, Medicaid, exchanges in the U.S., but also in every regulated, competitive health insurance market around the world.

Okay, so that's all just sort of setting the stage for where we start in this paper, and where we start is a couple years ago we started observing these kinds of reports in the papers that describe the idea that patients are being discriminated against in terms of the prices that they're paying for their prescription drug coverage. So I sort of pasted on a few of the headlines. "HIV Patients Excuse Health Plans of Using Drug Costs to Discriminate." "Health Insurers Discriminate Against Patients who Need Specialty Drugs." The idea there is that, you know, or at least in many of these stories or in the consumer complaints that were coming in through HHS was that, you know, even though prices couldn't be -- premium prices couldn't be differentiated across people with different health status, that somehow this was still working its way through to the benefit designs.

Now, when we saw this as economists, we thought, okay, one of two things is happening: either it's the case that insurers are still operating in this -- because these are markets with risk

220

adjustment, so either it's the case that insurers are still set in this mind-set that a costly patient is an unprofitable one, or they're actually correctly understanding the incentives, with some level of sophistication, and what they're finding are the places where the risk adjustment is sort of, you know, not properly calibrated in some sense. There is still some error, some margin -- some margin for profitable selection.

And so that's what we look at in this paper, and so there are these anecdotes pointing to the idea of limiting access to entire classes of drugs as a backdoor for discrimination, and the kind of complaints and the kind of statements that you would see HHS making, but also the complaints that are coming out of consumer groups, were that most or all of the drugs that treat some specific condition -- so, you know, the whole set of alternative substitute therapies -- were placed on the highest cost-sharing tier. So it's that anecdote that in the paper we're going to evaluate systematically, and data.

So what do we do in this paper? We're going to study this kind of selection-related formulary design, so the way that plans are creating their prescription drug formularies, using data from the 2015 ACA

221

exchanges, or now they call them marketplaces for as long as they still exist, and we investigate whether the drugs treating chronic conditions are, first of all, just trying to figure out, are they a plausible screen? Can insurers actually, at least in principle, make money by selecting -- by selecting consumers with this kind of screen.

And the reason why you might think that prescription drugs are the right place to look for this kind of activity is, you know, among all the kind of healthcare goods that healthcare consumers consume, you might think of drugs as being -- especially drugs that treat chronic conditions where I need to take this drug every month -- as being particularly transparent in terms of both need and price.

So we're going to sort of, you know, in the next 20 minutes ask and answer two questions. First, is there scope for selection? So is it the case that there is some problem with the risk adjustment system that's leading to the ability to profitably screen certain kinds of consumers? The answer there, I'll show you, is yes. The second question is, you know, saying that that incentive exists is one thing, and then the question is, are there -- are the insurers appearing to respond to that? And the answer there is

222

yes, and with, to my mind, what's a pretty significant level of sophistication. I'll tell you, as we go, sort of why we think that's the case.

So just to give a little bit of orientation on the literature, and I won't spend much time here, you know, the first talk that we heard today was based on the Akerlof lemons model. In health insurance, the way that we apply the Akerlof lemons model is we think about selection impacting the composition of a risk pool and then ultimately feeding back into prices in a competitive or imperfectly competitive market, and there's been a lot of both good theory and good empirical work on that, Einav, Finkelstein, and colleagues.

One thing that that model really can't say anything about, because it assumes it away, is the kind of phenomena which you didn't hear, which is that the contract itself changes. It's not just that you change the risk pool, and by changing the risk pool, you change the break-even price in a competitive market, but it's that insurers are not sort of passive participants. They design plans, and they can design plans with these ideas in mind. Of course, this is kind of the original idea of Rothschild Stiglitz, but there's also been other good empirical or theoretical

223

work thinking about this idea applied to health insurance markets. Where there's a gap is that there's almost no empirical work on this.

So in this paper there's basically no theory. We're just taking sort of the envelope of insights and kind of empirical predictions from the existing theoretical literature -- so Veiga and Weyl, Azevedo and Gottlieb, some papers by Tom Maguire, and of course Rothschild and Stiglitz, and we're going to take that in and we're going to look for empirical evidence.

Okay, so the first part of the exercise is just trying to understand, you know, how well is the risk adjustment working? Is there plausible space to use formularies as a way to screen out unprofitable consumers?

So I will try not to make you learn more about healthcare regulation than you absolutely need to to get through the slides with me, but there's two broad categories of regulations that are intended to deal with this problem. The first are things like a coverage mandate. So in the Affordable Care Act, in the exchanges, there are things like essential health benefits. This is where the regulator says to the insurer, you must cover X.

224

The other family of regulations are payment adjustments. So rather than saying you must cover X, even though X is going to attract unprofitable people to your plan, instead what we will do is we will adjust the payments so that those facts, on net, after the risk adjustment or reinsurance, are no longer unprofitable.

And so I've mentioned a bit about how risk adjustment works, but what risk adjustment is doing is it's going to make a payment to an insurer based on the diagnoses and demographics of the people of the risk pool that's enrolled in its plan. And reinsurance is going to make a payment based on the ex post realized healthcare costs of people enrolled in the plan.

Okay, so to answer the first question, which is about do these incentives exist net of risk adjustment and reinsurance, we're going to go to detailed health claims data. These data are not going to be from the marketplaces, the exchanges themselves. It's going to be out of sample, because that's where we can get claims data. And what we're going to do, in those claims data, we will see a person's costs, and we'll ask the question, what would the risk adjustment and reinsurance payments have been if this person generated that claims history while enrolled in an exchange plan?

So, you know, we can just take off the shelf the algorithm from the regulator, HHS, and we can say, here's what the risk adjustment payment would have been, here's what the reinsurance payment would have been, and ultimately here's how unprofitable or profitable this consumer would have been if enrolled in your plan and generating these claims in an exchange plan.

So just to give a bit more detail on how we do this, premiums here are not sort of, you know, completely stable or completely constant across all people. We're just going to take the average cost in the sample and assign it actually a fair premium. All the variation that we are going to be identified off of is the implied risk adjustment, and implied reinsurance risk adjustment, remember, is a function of diagnoses and demographics, and reinsurance is just a function of did you -- did you generate claims that were in excess of some attachment point, at which point the reinsurance kicks in?

This gives us profitability at the individual level, and then what we want to do now is try to connect to the anecdotes that said what insurers

appear to be doing is taking all of the drugs that treat some condition and moving those to a restrictive tier, and in the complaints, usually the specialty tier of drugs. So we are going to group consumers within therapeutic classes of drugs.

So some examples -- so we are just going to take a standard issued definition of these classes. So anticoagulants or blood thinners or statins or oral contraceptives, antidiabetic agents, these kinds of classes within which we think there are substitutes or alternative drug treatments, and we're just going to -- we're going to take all the folks that use one of those drugs, we're going to calculate the average cost, conditional on a flag for that -- on a drug for that class, and look at the expected revenue conditional on that same flag.

And what comes out of that -- I'll try to do most of this in sort of nonparametric plots. What comes out of that is a scatter plot that looks like this, and so the -- each circle is a different therapeutic class of drugs. The position on the horizontal axis is the average cost of people who use a drug in that class, and on the vertical axis, the average revenue of people who use a drug in that class, and the size of the bubble is proportional to

the number of consumers in our data that use each class.

What you see is a couple things. The first fact -- empirical fact that comes out of the analysis is that for most classes, the selection incentives are pretty well neutralized. So, a 45-degree line tells us that, you know, even though someone that takes a vasodilating agent to treat chest pain is going to have $4,000 in expected costs, and the insurer knows that in some sense at the time that the person is enrolling, if they knew that they wanted that drug, they're also going to generate about that amount in revenue, because there's a small premium, but there's also an $18,000 risk adjustment payment for someone who shows up with that diagnosis in your risk pool, and there's another $4,000 or so in reinsurance payments, right?

So for most consumers in most drug classes, these incentives are really well balanced, and I think that's pretty interesting and not at all a necessary outcome since the risk adjustment algorithm doesn't actually take into account what drugs you take. It takes into account your diagnoses, and that's going to be correlated to some degree with the drugs that you take.

But it's not universally true, so there are these outliers. So, for example, by logical response modifiers treat multiple sclerosis. A person who's going to demand a drug in this class is going to generate an expectation of $61,000 in costs but much less in revenue, even after taking into account $34,000 in risk adjustment transfers and a sizeable reinsurance payment of, like, $9,000 as well.

And so when we go on to the second part of the analysis, we try to see, you know, are plans responding to this incentive? What we'll look at is basically vertical deviations from this 45-degree line, and that vertical deviation is, you know, in dollar terms, in level terms, how unprofitable is a person who predictably will demand a drug in this class?

Very briefly, something else that came out of this which I have to mention because it -- unless I mention it, I don't think it will come across just by looking at this last picture, is that there's absolutely no correlation after risk adjustment and reinsurance between costs and profitability, and what that's going to mean is that when we look at insurer sophistication response to this, the insurer has to be more sophisticated than merely saying we are going to

229

try to keep expensive people out of our plan, because there's no longer any correlation between expensive people and unprofitable people.

Just for time, let me skip this, just a different look at the people that you want to avoid. So why are there errors in the payment system? So, you know, one possibility is that in the time between the payment system being calibrated and used, there was some technological change in how -- you know, what -- how costly it was to treat a particular condition, but, you know, more generally, there's no reason to think that these things would be orthogonal to profitability since they weren't included in the -- in the algorithm that tried to -- that tried to net profitability to zero for each group.

We will skip that for time. All right, so then the second goal of the paper is trying to ask, you know, not just do these incentives exist, but do plans respond to them and with what degree of apparent sophistication. So for here now we'll actually go -- so all of that so far has been a sort of out-of-sample prediction, looking at basically employer health plans, large self-insured employer health plans and the claims generated there. So now we want to ask, do drugs that predict unprofitable patients, are they

230

actually covered ungenerously by exchange plans? And so for that we will turn to exchange data. So we will get the universe of formulary data from 2015 from the exchanges. We'll do that both for -- so we will look at the exchange data, and we will also use employer plan formulary data as a sort of comparison point, so we can do a difference in differences. And the unit of analysis here is always going to be at the drug class, so grouping together all the potential drug therapy substitutes by plan.

When -- as we go forward, as I keep talking about restrictiveness, what I mean by restrictiveness is if you took a plan's cost-sharing tiers and you sort of ranked them from most generous to least generous, there's a very clear breaking point at the level of specialty drugs, and one of the reasons for that -- although we could talk more about this if you're interested -- is that's generally a level at which you go from a copay regime, so you pay 30 or 60 or 90 dollars, whatever your plan says, to you pay a coinsurance rate, 20 percent, 25 percent, whatever it is, and that could be a really important difference, and we show that in the paper for high-cost drugs. It's also the level at which there's -- you know, states have taken regulatory action. So, for

231

instance, New York State outlawed the use of a specialty tier in health plans in the state.

So when I say restricted drugs, I'm talking about specialty drugs or drugs that are left off of the formulary altogether or drugs for which, if you want to use them, there needs to be some sort of nonprice hurdles that you need to jump over, like step therapy or prior authorization.

So just very briefly, the -- we're going to be comparing in some sense employer plans to exchange plans and how they differentially respond to this selection incentive. The selection incentive doesn't exist in employer plans. They're just sort of a useful control group for us. Because it doesn't exist in these plans, they're not subject to the risk adjustment and reinsurance rules, but those plans are -- exchange plans and employer-provided plans, they're just differentially generous, so as we go forward, we will be controlling for that differential generosity.

So here's -- here's kind of the main result in a picture, although, you know, I'll show something with a little bit more detail in a minute. So what we're doing here in the -- on the left-hand side, we're taking the drug classes in the bottom tenth

232

percentile -- bottom ten percentile of these selection incentives. So these are the guys that are actually relatively profitable, where the risk adjust arrow is going in the direction of you'd want to -- you'd want to get these guys into your plan, and the two bars on the right are the 90th percentiles and up of this selection incentive. So these are the guys that you want to avoid. These are guys that demand drugs that you -- that predict unprofitability.

What you see is that there's really -- there's no gradient here in employer plans, nor should there be, because employer plans aren't subject to these incentives, but we want to -- we want to use employer plans as a control group, because we want to control for the fact that, you know, some drug class versus another drug class might be more subject to moral hazard, where it might make sense to be -- where it might just be more expensive. There might be sort of good, efficient reasons to restrict consumer access to these drugs or make it a little bit harder, but what you see is that across -- across these percentiles of how profitable or unprofitable the patients are, there's basically no reaction in the employer plans and a relatively large reaction as a share of the drugs that are restricted access in the exchange plans.

So here's taking that kind of difference-in-difference idea and just doing it a little bit more fleshed out and fully. So here what we're doing is we're taking all of the drug classes and we're grouping them into ventile bins, so there's 20 points, and within each ventile bin, we're saying, all the way on the left, these are the drugs that point here (indicating), these are the ten classes of drugs that are the most profitable, and we're asking along the vertical axis, how frequently are drugs in that class restricted access in exchange plans relative to employer plans?

All the way to the right are the least attractive drugs or the drugs that predict the least profitable people. And so, you know, one of the things you see here is that most of the points are in the middle, sort of the neutral selection incentive, but that's -- that makes sense because most of the points lie along the 45-degree line in the first picture I showed you, so it's only really in the points where we -- in these sort of outlier points where sort of the entrant binds in some sense, where there's a mistake that then we can see how insurers respond to that payment system mistake.

You know, I don't think -- in terms of the coefficient estimates here, I don't think there's anything more important to glean than what you can see on the nonparametric plots, but just so you can understand what we do as we go forward in some additional specifications, the actual regression that we run in the parametric specifications are, we have drug class fixed effects, statins, anticoagulants, what have you. We have plan fixed effects. Then we're asking how, within a plan, across drug classes, does this selection incentive predict how generously or ungenerously, relative to other drugs in the same plan, the drug is covered.

And so we find that -- we find that these drugs that predict unprofitable people are both tiered ungenerously in terms of being on specialty tiers more often. They're also tiered less generally in terms of -- generously in terms of requiring some kind of nonprice hurdle to be met. So whether it's step therapy, we have to try other drugs first, or a prior authorization, where you need to call the insurance company, and there are reasons why we think that might be important in this context, which we talk about in the paper, having to do with the cost-sharing subsidy reductions.

Okay, so with the last four minutes, let me say a bit about insurer sophistication, because, you know, while I think it's good to have these kind of parameter estimates that come out of the paper, to me, the story of the paper is do insurers respond to these incentives, yes, and how sophisticated are they in responding to these. That's where I think we have some interesting things to say.

So what's important is in this setting, the drugs themselves are a small fraction of cost. So here we're using the same ventile bins from the most profitable group all the way to the left to the least profitable -- most unprofitable group all the way to the right. Those are the guys you want to avoid. In all of these cases, drugs are a relatively small share of the costs. So the drug is a signal for the patient profitability. It's not actually the thing that's driving that profitability or unprofitability. And as I showed before, there is no correlation in overall cost in patient profitability. So there has to be some level of savviness on the part of insurers if they're actually responding to these net of risk adjustment and reinsurance incentives.

So we spent a lot of time thinking about this in the paper, because this is something we really want

to understand, and one of the things we do is we start just dividing up this graph into vertical slices, where we're looking at just patients that are equally costly but differentially profitable.

So just relaxing the parametric assumptions even further, looking within vertical slices, so folks that take cardiac glycosides, vasodilating agents, and gonadotropins, these are all people who are going to generate the same healthcare costs, roughly, in expectation, but they're also, in expectation, going to generate very different profits. And the fact that insurers are responding to that profit motive indicates some, in our minds, serious sophistication.

So with two minutes left, you know, these are just the regression specifications that show that kind of comparison within vertical slices. I don't think it's useful to go over them, other than to say that we get the same results when we condition on these vertical slices, meaning patients with the same underlying expected cost.

Also, just to summarize, in the paper, we do a lot of work ruling out other alternatives, potential explanations for this, you know, like is it the case that this selection incentive is correlated with moral hazard? You know, some drugs, if -- if there's more price sensitivity within a drug class, then you might want to restrict access to it in the same way that -- well, I'll leave that there.

Is it just about nudging consumers towards more cost-effective options or generics? No. Even when we look just at the generics within a class, the generics are more likely to be left off of a formulary if they predict a patient is going to be unprofitable, if they have predicted an unprofitable enrollee. So it's not about nudging to cost-effective alternatives. It's not about a nudge to go generic. It's also not about nudging consumers to products for which the pharmacy benefits manager gets a better deal.

So we can do all of this by looking within pharmacy benefits managers and saying, here are -- here's United Health Plan's employer plans in Texas, and they use some PBM. Here's UnitedHealthcare's exchange plans in Texas. The same insurer, the same PBM, generates very different formulary structures in the two markets and in a way that's correlated with the selection incentive that we document.

Okay. So with the last half minute, just to conclude, some important take-aways here are, first, even though what we're interested in here are the deviations where there's a breakdown in the risk

adjustment system, overall, it's important to understand that risk adjustment and reinsurance are doing a pretty good job of protecting consumers with preexisting conditions from having plan designs tailored against them. So here we're looking at drugs, but, you know, you might think about hospital networks being formed in a way, you know, leave out the really good cancer hospital if cancer patients are unprofitable, and risk adjustment and reinsurance seem to be doing a good job in this sense overall.

Where we see deviations, where you can predictably tell who's going to be profitable based on the drugs they demand, we see insurers following those incentives. And it's not about high cost. It's -- insurers are sophisticated enough to understand who's profitable.

Then a couple last notes on regulation, I think a lot of the ways that policymakers and regulators often think about this is what we need is really strong essential health benefits controls, where we need to say that you must -- you must, you know, cover some drug in each class. Those are in place in the Affordable Care Act health insurance exchanges.

The problem is that this product is incredibly multidimensional, and there's just -- there's no way

to regulate all the dimensions along which insurers can design their plan to try to cream-skim enrollees. If you say that, you know, you can't put these drugs on the specialty tier, then they will put more nonprice hurdles in consumers' ways.

Really, the only way, in our view and from this paper, to ensure this kind of access is to get the payments right to the insurers to remove the financial incentive. I'll leave it there.

(Applause.)

MR. ROSENBAUM: Discussing the paper is Sebastian Fleitas.

MR. FLEITAS: Okay. So thank you very much, and thank you (off mic). Oh, sorry, yeah. So this -- okay, there we go.

Okay, so the idea here is that risk adjustment and reinsurance introduced in the exchanges is a way to compensate for enrolling costly employees, so this is important why? Because we don't want to deny access to these -- to these enrollees, and basically we wouldn't want to price-discriminate them, because then they will be exposed to risk -- reclassification risk, so we don't want that. So basically we want these schemes to work and to -- and to make equally profitable to enroll up a consumer that is very

costly, okay?

What this problem is, the problem is that this mechanism may not work well, so maybe we can have issues with this. And basically on top of this, I mean, the firms can try to make actions, try to do things in order to screen out consumers, okay? So if I understand that these consumers are very unprofitable, I try to design my formulary, for example, in some sense trying to get this consumer out of my plan, okay?

That's a screening mechanism, and the existence of the extent of this screening mechanism is actually an empirical question, okay? We want to see in the data what's going on, okay?

So this paper basically is going to do two things, as Michael said. So the first thing, it's going to show that actually with this (indiscernible), that the adjustment and the insurance works pretty well for a lot of drugs, for a lot of drug classes, so in that sense, it's pretty (indiscernible).

But also the paper shows that there are some payment errors, and these payment errors can be used for a screen, and actually, this is trying to show, okay, what's the strategy there? The strategy is to use a difference-in-difference approach, okay, and

241

basically it's going to compare the exchanges, so what happens in the ACA, with the employer-sponsored insurance.

Basically the idea is that this way we can -- so as Michael said, there is no incentives to deviate with these mechanisms in the employer health insurance, so basically the idea is that the difference in differences is going to allow us to control, by plan and drug class, fixed effects, okay?

So we want to control whatever is the same for all the drug classes and whatever is the same for all plans, okay, and then we want to use a gradient of the drugs to identify the model, okay?

So basically what we're assuming here is part of the trends in the class-specific costs and revenues, okay? So this is the main assumption of the paper, and we are going to go through these in a sec, okay?

So basically let me tell you that this is a -- as you may see by the presentation by Mike, basically this is a very nice paper. I think it's important. It's actually transparent. The paper is very detailed, so it has a lot of results, so we can track what's going on here, and it's very clear. So basically it was a pleasure to read the paper.

242

I think it's an important paper. So basically, as I said, two main results. So the main thing, (indiscernible), these things works relatively well. The second thing, there are still errors. The firms are using these errors to screen consumers. It is important, obviously, for policy reasons.

And actually, the second thing is that here the insurers are relatively sophisticated, okay? So they can understand that cost is not the same as revenues minus costs, so they can do that, and this is important in the way we design the mechanism, okay? It's very important for the policy, okay?

The main contributions, basically this paper adds to the literature that highlights the important role of nonprice characteristics in strategic behavior, and this I think is important for three things. First, to understand the use of screening strategies by firm, so generally in -- in how they work. For regulation, it's extremely important, because we need to understand how much these remedies can alleviate the problem. So, for example, essential health benefits or the risk adjustment system, how to compute those.

And obviously, for modeling in economics, it's extremely important, because it makes characteristics

243

not being exogenous, okay? So it is extremely important for us, because then in the dynamic setting, for example, we introduce the endogeneity of the characteristics, and it is extremely problematic, for example, if we have a setting which is multidimensional, when we have a lot of state variables, for example, the (indiscernible), okay? So this is going to be a problem. Also it's important because we have the ACA, which is a relatively new and important market, okay?

So let me tell you very briefly three comments about this paper that I have, maybe some things that are there, so...

The first thing is that I see the paper, we see a lot of the evidence, so we compute the cost -- the average cost and the average revenues, okay, but we don't play that much with the standard deviation, okay? So since you are having to use a lot of data, you can compute actually what's the standard deviation here.

This can be important because we would like to understand if this -- if this standard deviation is coming from consumer heterogeneity or it's coming from cost heterogeneity. So it's treating different conditions, okay?

244

So maybe one thing to do is just to use the standard deviation of this cost minus revenue measure in the regression, you know, interactive with exchanges, obviously, and trying to see if this gradient also change with the standard deviation, okay, if there is some response with this, okay?

But a little bit more problematic is that this opens -- if consumers are heterogenous, this opens some challenge to identification, and basically the idea is that the selection mechanism works basically pricing out of the market some consumers, okay? So placing restrictions to leave some part of the consumers out of the market, okay?

Therefore, this may lead to different elasticities of consumers who are in the two markets, and by the two markets, I mean the exchanges and the employers, okay? So basically the price elasticity of these two guys, of these two people are going to be different, and basically the drug and plan fixed effects don't account for this, okay, because these are specific to the plan under the right class, okay?

So basically here is -- the scope of selection here can be problematic with the heterogeneity of consumers, okay? And basically (indiscernible) is the main (indiscernible) assumption, okay? This also can

245

be problematic in the sense that this selection, these guys that don't -- are not being covered in the exchanges can go to other places, so maybe they go to exchange insurance.

The employee insurance is also, like, changing the composition of this group, so maybe the cost for one part of this group is not the same as the cost of the other group, okay? So in that sense, obviously, maybe the small group market is a more close substitute than the ACA, but in any case, I mean, it would be nice to see what's the flows of these guys going from one market to the other, because this can generate some issue, okay?

So I think one thing to -- we can do to try to understand this a little bit, obviously higher variance of price elasticity is going to open more opportunity for reselection, okay, for having different types of consumers in the two different markets. So one way to do this is to estimate this different price elasticity in these (indiscernible), okay, and use this -- this -- again, the amount of price heterogeneity of the -- of the elasticity -- of the heterogeneity and elasticities in order to also use with exchange, okay?

The idea, as I said, is just if you have more

246

heterogeneity in one class than in the other, it opens more opportunity for having more selection in one or the other, okay?

So the second comment is also -- is a question, like how -- if there is an also story about higher costs, okay? Not in the way of sophistication that we discuss in the presentation, but in a different way, okay? So basically the idea here is that the firms are using these formularies in order to send a signal for consumers, saying if you are unprofitable consumer, don't come here. This is seen as an (indiscernible), because you are costly for me, and I don't want you here, okay?

But also it can be -- so it can be that they respond this way because they have a higher cost with these particular classes, okay? There is going to be also a story of costs or a story of pass-through, that these firms are actually sending a signal saying, okay, we have a higher cost using these drugs, so we will send it to you, okay?

So the first thing to note is here is that if we see a table in Appendix A-1, okay, so with the three measures of profitability, and we see the number on (indiscernible), we see that most of the results concentrate in the last ventile, okay, at least in the

247

three last ventiles, okay? So this is something that is happening in the class that are very unprofitable, okay? It's not having a (indiscernible) distribution, but it's mostly having the ones that are very unprofitable, okay?

And the same is true when we control a pharmacy benefit managers, okay? This is Table A-9 in the appendix, and we can see basically the effect seems to be very, very concentrated in the ventile number 20, okay? So this is something that happens with this class, okay?

And then you can see again the same graph that actually Michael presented in the presentation, okay? So it is true, as Michael said, that all drugs represent a small fraction of cost all over the different classes, okay, but it's also true that if you see how the share of drugs represented here increase relatively high in the upper -- up from ventile number 16, actually, okay? So it's very correlated, the share that the drugs represent in terms of cost, okay, with the very unprofitable condition.

I think that the clear reason for that is that drugs utilization is not using the risk adjustment mechanism, okay, so in that sense, this induces

248

correlation, okay?

But maybe it's just the firm is saying, okay, I have these -- I mean, these costs are much higher, okay, they use drugs, so one way I want to reduce this or tries to do some pass-through to consumers is to increase the cost of these drugs, okay?

So I think one thing -- one easy thing to do here is just to control for the -- for the share of drugs that comes here, okay? So basically using the same regulation, just use the share. In that sense, what you want to use as a defined variation is the changes in other expenditures and nondrug expenditures, okay?

So this is very easy to do, so easy, but it can be informative of how much of that pass-through is a story -- a cost story and how much is a screening story, is that when drugs not are even in my cost, okay?

So the third thing and last thing is about competition in this market, okay? So basically obviously with these incentives, the firms are going to respond in equilibrium. So the first thing we would like to know is how much heterogeneity we have by market characteristics. This is also kind of easy to do.

So just interact the exchanges with some market characteristics, like the number of competitors or these kind of things, can give you some idea of how this -- how they respond. So this also, like, can give you some -- some clues on why we see all the effect concentrating on the very unprofitable consumers and not a lot of fight -- or maybe a lot of fight, actually -- with the profitable consumers.

And maybe they are competing very, very hard with the profitable consumers, and that's why we don't see effect there, and they try to get all -- everyone is trying to get rid of the unprofitable consumers, okay?

Also, the last thing is our dynamics here, okay? I think there are at least two sources of dynamics that can be interesting to understand here in this market. The first one is the learning, okay? So basically we have a cross-section of data, so we are going to go through that, but learning here can be important. I mean, obviously, the firms need to learn how to play this regulation, how the adjustments are doing, so basically learning is an important perspective here.

The second one is inertia, okay? So basically the inertia has also been documented in healthcare

markets, and we have a dynamic competition. There are some elements of dynamic competition, some markets, for example, in Medicare Part D, so we see clearly these investing and harvesting dynamics, when the firms further reduce prices to capture consumers, and then increase the prices to exploit them. Something like that can be happening here, okay? So there are ways to do this.

So basically the easy way that doesn't require any extra information that you have available is using the vintage of plans in the market, okay? So basically you have for the exchanges every plan that was offered in each of these -- of these markets, so you can check what's the vintage of this market, what's the number of years a plan is in the market, okay, and use this variable to check that, okay? So you will inspect, like, some effects on the industry.

Other ways to do it is just using market shares of plans by condition, trying to see if the plans -- the newer plans have higher market share -- lower market share of -- of profitable conditions and higher market shares of unprofitable conditions, if that makes any sense, or maybe approximating the market share by condition using shares of expenditures. Maybe that's easier way to do it, but it's true that

it's difficult to get information -- individual information for exchanges.

Okay, that's it. Thank you very much.

(Applause.)

MR. ROSENBAUM: We have time for one or two questions.

MS. JIN: Thank you.

I think the results are really fascinating. I have two questions. One is, how does this relate to the market power of pharmaceutical companies? Is it true that unprofitable drug classes are the ones that has more market power on the pharma side, and that's why they charge super high price to the insurers and sort of force the insurers to either sort of drop the drug or use other tactics to deal with the high cost?

Another question is, what's the consequence of this? Is this just like every exchange plans refuse to offer coverage for that kind of drug so that it completely shut out this market to that type of patients or it's just more differentiation story in the sense that there's still at least one plan offering -- offering this kind of drug coverage? It's just not as many plans as in other classes, so that could sort of consolidate all the patients in exchange market into that particular plan, and that could be a

differentiation story that makes this plan more differentiated from other plans.

MR. GERUSO: Yeah, thank you.

So to your first question, the -- you know, I don't think we -- this paper sheds any light on the role of market power of the -- of the pharma manufacturer. In part, that's because the way this -- the whole exercise is structured is that we're going to difference out anything that's constant within a drug or within a drug class, because we're comparing how these drugs are tiered in employer-sponsored insurance plans versus the exchange plans, and that's -- you know, that's an intended feature, but that means that that -- you know, that pharma market power is going to be constant, at least in some sense, between those -- between those two insurance delivery mechanisms.

And then the idea about, you know, is this about differentiation, we tried to -- we tried to dig into this a little bit. We've got future work planned where I think we'll be able to get at this a little bit better and also get at some of the competition questions that Sebastian was bringing up, but, you know, our initial cuts of the data where we just looked at let's just divide this kind of -- you know, there aren't that many insurers, right, so let's divide this by big carriers, and let's also look, like, at small, non-national carriers, and we just didn't see -- although we're limited in our statistical power to detect -- we didn't see differences across carriers. That's not exactly the same as the question of, you know, is there some carrier that's left offering this plan, but that's something we are digging into in the next project.

I mean, I will say that one of the facts that motivated this was that this paper in the New England Journal of Medicine by Jacobs and Sommers that was pointing out that in Florida basically it was impossible to get a plan that covered HIV medications on less than a specialty tier, sort of regardless of what the actual underlying cost of those medicines were.

So I think it's possible that we're in a symmetric equilibrium in which, you know, no plan wants to be the plan that's left holding the bag with the unprofitable patient, but, you know, certainly theory -- there's a lot of theory and very little evidence in this area so far, and there are -- you know, some of the theories are symmetric equilibrium, some about a -- you know, it was more separate in

Rothschild Stiglitz. So it's an open question, but we hope to work on it.

MS. SAEEDI: So I have a question. If the cost that you're showing is only the drug cost or also the cost of everything else that is the patient cost and if they are correlated and if they are trying to sort of avoid these customers from just getting -- sign up just -- not directly because of the drug cost but because they cost higher...

MR. GERUSO: Yeah. So, you know, on average, drug costs are 20 percent of the patient cost or even less -- they're among -- conditional on having any drug use, it's something like 20 percent, and even among the most unprofitable patients only rise to 40 percent, and we're doing -- in doing our profitability calculations, we're using all costs -- hospital, inpatient, outpatient, visits with doctors -- using all of that to figure out profitability, because that's what matters to the insurer.

To connect what you're asking with a comment of Sebastian, when we -- when we start trying to figure out what are insurers responding to, I didn't really -- I skipped over it, but we show that they are responding to costs, they are responding to drug-specific costs, but even netting those things

out, so just controlling for them simultaneously, they're simultaneously reacting to the profit incentive sort of conditional on the costs.

I think that's something I didn't get across clearly in the original presentation, but, you know, we can control very nonparametrically, very flexibly for cost. You know, drug costs, nondrug costs, we can put that in there. They're responding to that which, you know, is interesting, but they're also separately and without sort of diminishment responding to the profit incentive.

MR. ROSENBAUM: Okay. Our next presentation is Fernando Luco, presenting Multiproduct Firms: When Eliminating Double Marginalization May Hurt Customers.

MR. LUCO: Perfect. Thank you.

Well, first, thank you for having me here. This paper is joint work with Guillermo Marshell from the University of Illinois. What we do in this paper is to think about markets that look very much as what Paolo was discussing -- talking about an hour ago.

What we want to do here is to think about bilateral oligopolies where upstream and downstream firms interact with each other and with consumers, and, in particular, we want to think about vertical integration in these markets, okay?

And often what we do when we think about vertical integration is to think about the trade-off between the efficiency gains that are specific to the transaction, eliminating double marginalization, and at the same time, market foreclosure (indiscernible), that may lead to increasing costs -- increasing the cost of doing business for some of the rivals, so the vertically integrated firm.

Now, this idea -- and, in particular, that we often assume to some extent that efficiencies are going to be realized -- has driven some very important economies to suggest that we should approve vertical integration unless there are clear incentives for foreclosure.

However, this is coming mostly from a literature on single-product firms, and what we do in the paper is just think about what happens when we talk about multiproduct firms. And, in particular, what's going to happen here is that to foreclosure and to efficiency gains, we are going to add a third effect that has to do with how partial elimination of double margins may lead to actual price increases that may hurt consumers even in the absence of market foreclosure, okay? So we may have just efficiency gains, and this third effect may actually lead to

257

price increases.

So to fix ideas, let me start with a super simple version of the industry that I just show. Here I have two upstream firms, U1 and U2, that will sell different products, substitute products, to a retailer at prices, Omega 1 and Omega 2, and the retailer will resell the product to consumers at prices P1 and P2. A super simple framework, okay?

And in this setting, what we want to think about is what happens if you, one, integrate with the retailer. So the way we've framed the question is, well, what's going to happen here? We're going to eliminate double margins for product one. That's going to lead to a decrease in the unit cost of product one, Omega 1 is going to decrease, and that will have two effects.

The first one is something that we are very familiar with, is that a decrease in Omega 1 will put a downward pressure on the price of product one. That product is cheaper to produce, so it's going to put downward pressure on the product -- on the price of that product, and that is the efficiency effect. That's what we think of when we are thinking about eliminating double margins in these type of industries.

258

The second effect is sort of the new thing here, is that together with making product one cheaper, you're actually making product one more profitable, so that may lead the retailer to increase the price of product two to divert demand to product one, okay? So let me repeat that, because this is a key part of the paper.

So we have a decrease in the unit cost of product one because of the elimination of double margins. That leads to a decrease in the price of product one, but because the prices [sic] are substitutes, the retailer has the incentives to increase the price of product two to divert demand to product one.

Now, we're not the first to suggest that this exists, so, in fact, the literature comes all the way from Edgeworth, and so this is a reprint in 1925, and the original paper is from the 1890s or something like that, where he was talking about taxation -- and product-specific taxation -- and when he suggested this path for product-specific taxes, someone replied that this is one of the horrible things that happens when math takes over economics, okay?

And that's sort of -- people didn't look at that that much. It was called the Edgeworth paradox,

259

and it was not until Michael Salinger, in 1991, brought the data into vertical integration. So for this reason, we actually call these the Edgeworth-Salinger effect, and I am going to refer to that for the rest of the paper.

So in this context, what we do in the paper is to ask whether the Edgeworth-Salinger effect is something that we should take into account when we're talking about vertical integration. So we're going to ask, what is the magnitude of the Edgeworth-Salinger effect? Should we consider these in merger evaluation?

And, in particular, as you saw in the previous slide, while efficiency gains seem to drive prices down, the Edgeworth-Salinger effect seems to drive prices up, so we call it that these two effects play with each other.

That is going to put us, of course, in a very rich -- in the context of a very rich literature on vertical integration, both in theory and empirical, but our work is more related to the literature on the Edgeworth paradox.

To answer the question, we are going to use data from the carbonated beverage industry in the United States. So let me spend 30 seconds telling you

260

about the industry, and then I will tell you why we care about it.

So in this industry, we have upstream firms, such as Coca-Cola Company and Pepsi and Dr. Pepper/Snapple Group, that sells syrup to bottlers that have exclusive territories, and these bottlers can -- some of them can actually interact with more than one upstream firm.

So what I mean by that is Coca-Cola bottlers can bottle for Dr. Pepper. Pepsi bottlers can bottle for Dr. Pepper. Coca-Cola bottlers and Pepsi bottlers cannot bottle for Coca-Cola and Pepsi, okay?

Now, why do we care about this industry, aside from the fact that it fits the picture I had at the beginning? Well, because in 2009 and 2010, both Pepsi and the Coca-Cola company vertically integrated with some of their bottlers in the U.S., and this is going to be very useful for a number of reasons.

First, they didn't integrate with everybody, so that generates variation in particular structure across the country, and in a subset of the areas served by the bottlers involved in the transactions, these bottlers actually had licenses to sell Dr. Pepper products, okay? So we are going to see areas where nothing happened, there was no vertical

261

integration. We are going to see areas where there is vertical integration, but the bottler didn't have the license to sell Dr. Pepper. And we are going to see areas where there is vertical integration, and the bottler has the license to sell Dr. Pepper, so we can actually identify the Edgeworth- Salinger effect.

A benefit of this case, in particular, is that we have no evidence of market foreclosure, and that is going to allow us to have cleaner identification of the Edgeworth-Salinger effect. This is basically -- because at the moment the transactions took place, when there was a change in the ownership of the bottlers, there were termination clauses in the contracts between the bottlers and Dr. Pepper that were triggered, and both Coca-Cola and Pepsi went and reacquired the licenses to continue selling Dr. Pepper.

So they decided to continue producing these products, and at the same time, while the FTC cleared the transactions, subject to a number of nonbehavioral remedies related to how information regarding Dr. Pepper could be used by the vertically integrated firm, but market foreclosure was never really a major presence.

So what is the data here? We have some really

262

novel data, some data that you know very well, so I am going to be very brief about it. So the part that you know very well is the IRI marketing data set, that we have weekly scanner data for the years 2007 to 2012 across a number of regions in the U.S. Our observation here is going to be a store-week-brand-size combination, and we are going to focus on brands that have at least 0.5 percent of the market. So that's going to leave us with 105 products that will -- and that's, for example, a 67-ounce bottle of diet Coke sold in a particular store in a particular week.

Now, where are things going to get novel? We have an industry publication called Beverage Digest that produces maps of the U.S. with the territories of each of the bottlers for both Coca-Cola and Pepsi, okay? So think of these as you have a map of the U.S. with state boundaries, forget about the boundaries, and you put the territories of the bottlers.

From there, we are going to be able to identify which areas were affected by vertical integration, and we're going to intersect that, if you want, with FTC documents that identify, in the areas where these bottlers -- where the bottlers involved in the transactions were producing, in which of those areas

263

they were actually also serving Dr. Pepper, producing Dr. Pepper products. So in the end what we have is, like, three maps that we are overlapping to pin down in which of these areas each of the effects takes place.

So let me show you the data. Here I have maps of two parts of the U.S. The map on your right -- your right -- no, the map on your left is the northeastern United States, and the map on the right is the Houston MSA, and as you can see, it's color-coded.

So blue areas are areas where nothing happened. There was no vertical integration in those areas. Green areas are areas where there was vertical integration, but the bottler did not have the right to sell Dr. Pepper. And orange areas are areas where there was vertical integration and the bottler did have the right to sell Dr. Pepper. So that means that we can use, with the in-store product variation, cost of vertical integration to identify both the efficiency gains associated to vertical integration and the Edgeworth-Salinger effect.

The Houston MSA is useful to illustrate two things. First of all, as you can see, the whole MSA is treated in the sense that there was vertical

264

integration affecting the whole MSA, but only one of the counties in the MSA actually experienced the Edgeworth-Salinger effect. So this is going to sort of define at what level we are going to be defining treatment, okay?

And later, I am going to actually reduce the sample and we will do this some sample analysis using neighboring counties that were differentially affected by treatment, and as you can see, the Houston MSA is a good example of that.

Okay. So, of course, this means that what I'm doing here is I'm going to follow a difference-in-difference research design exploiting this variation of -- the within-store product price variation, cost with vertical integration, and together with that, there are a number of identification threats that we have to take into account.

Some of them, for instance, are what happens if the Coca-Cola Company, at the point -- at the time of the transactions, also changes the way it does advertising or it changes its (indiscernible) policy or things like that. So concerns like that we can address using finite structure of our data.

Other concerns that are more directly related to the research design has to do with differential preference, for example, and we -- in the paper, we explore those using both summary statistics, and I am going to show you later a dynamic difference-in-difference version of our estimation that basically shows that we don't have differential preference, okay?

So let me jump directly into the results. So the most -- I'm sorry, no, I forgot one. Let me show you the equation first. So we're going to study how vertical integration affects prices, and we are going to divide these in two parts. So first we're going to see how vertical integration affects prices of Coca-Cola and Pepsi products when these are bottled by a vertically integrated bottler. That is, we want to estimate the efficiency effect of vertical integration.

So in this estimation -- in this equation, that's going to be captured by the coefficient B-own, that we're referring to owned brands, as the brands owned by Coca-Cola and Pepsi, when these brands are bottled by other vertically integrated bottler.

We are also going to distinguish the effect of vertical integration on Dr. Pepper brands that we call the Edgeworth-Salinger effect, and to do that we're

going to have this B-Dr. Pepper coefficient that's associated with Dr. Pepper products that are bottled by other vertically integrated bottler.

And then in the third line we have a rich set of fixed effects that is meant to address the identification concerns that we have in mind. Some of them, for instance, firm with fixed effects, for example, are going to allow us to tackle changes that may happen at the parent firm level. Then we are going to have county with fixed effects to address local shocks. We are going to have store and county product seasonal fixed effects, that they can take into account seasonal effects and local conditions.

And the other thing, we are going to use that treatment at the county level to class our standard errors, but, of course, we have done a lot of robustness in all estimations, okay, and the results don't really change.

So now, yes, let me go into the results of the paper. So this is the most important table, so everything that comes later is digging deeper into what is going on here. So if you want to keep one result in mind, keep this one. So I have two coefficients here. The first coefficient is the average effect of vertical integration on own brands;

that is, the average effect of vertical integration on Coca-Cola and Pepsi products when these products are bottled by a vertically integrated bottler.

What we have here is the prices of these products decreased 1.4 percent, so this is a manifestation of the impact of efficiency gains on prices. So this is the effect of eliminating double margins for Coca-Cola and Pepsi products.

At the same time, we have that prices of Dr. Pepper products went up by 3.9 percent when bottled by a vertically integrated bottler. This is the Edgeworth-Salinger effect, okay? So the price of -- the price of Dr. Pepper products are going up by almost 4 percent, and this is consistent with what Edgeworth brought out and what Salinger brought to vertical integration.

Now, I note that Andrew is going to bring out these later, so there is a back-of-the-envelope calculation here. If you weight these coefficients by premerger market shares, we still get that the average price paid decreased by 0.9 percent, okay? So I'm not saying that this is -- like, the merger is not welfare-increasing or anything like that. I'm just saying the Edgeworth-Salinger effect is relevant. It's huge. It's the same order of magnitude as the

efficiency gains, and it definitely has an impact on prices, okay?

If you look at what happens with -- what happened with listed prices, not paid prices, we get the prices increase on average by 1.8 percent, but what I -- but the table that I like -- I really like is this one, where we allow all these coefficients to vary by parent firm. So we estimate different effects for Coca-Cola and Pepsi.

And what you see here is a number of things. First, prices of Coca-Cola products and Pepsi products went down by 1 and 2.1 percent following vertical integration, and prices of Dr. Pepper brands went up by 3.1 and 4.2 percent following vertical integration. So both firms or bottlers of both firms reacted in the same way. So the effects are going in the same direction because they are basically reacting to the same incentives, the same changes in incentives.

One could be tempted -- and I was tempted -- to say that the firm that had the largest Edgeworth-Salinger effect, Coca-Cola, for 0.2 percent, also had the smallest efficiency effect that would be a consequence, for instance, of price complementarities; however, we cannot reject the equality of the 1 percent and the 2.1 percent. We can reject the

269

equality of the 4.2 and 3.1 percent, okay?

So the take-away from this slide is both firms are reacting in the same way to the changes in pricing incentives that are caused by partial elimination of double margins.

When you look at this over time, so this is the dynamic difference-in-difference estimation, you see other things. First, why do we need this? For two reasons. First, we want to address the question of whether or not there are differential preference. And as you can see, that's not the case in the preperiod, but second, you want to see when the effects started to take place, and what we see here is that the effects started to take place after the transactions, and it particularly was -- the effect was long-lasting. So I have no idea what happened with the second-to-last point over there, but basically we have very persistent effects over time.

Then -- so Ali suggested to these a couple of weeks ago, we repeated the analysis at the product level, okay, so here what we're doing is we're estimating one coefficient for each of the products in the sample that at some point, somewhere, were affected by vertical integration. If the story of the Edgeworth-Salinger effect is true, then we should see

270

that Dr. Pepper products have their distribution of coefficients to the right of zero, because we're expecting prices of Dr. Pepper brands to increase, and that is precisely what we see here with two exceptions.

With own brands, things are a bit trickier, because we say, well, the efficiency effect is going to drive those prices down, but the Edgeworth-Salinger effect may actually drive those prices up once you take into account price complementarities, and what we see here is that it's like half and half. So some of the owned brands have price decreases; some of the owned brands have price increases.

We can re-estimate these for quantity regressions, and when you limit the analysis to products that in this regression have significant coefficients on vertical integration, we get elasticities between -- the minimum elasticities are between minus one and minus 3, which is in line with what people have found before in this literature.

So let me spend -- sorry, no, I forgot this. So this is one of my favorites. One of -- okay, so what we do here is to do a subsample analysis where we drop all counties that were exposed to the Edgeworth-Salinger effect. So we repeat the analysis

271

looking only at counties that either did not experience vertical integration or experienced vertical integration but didn't have the Edgeworth-Salinger effect.

So what we want to do is to compare the efficiency gains of vertical integration to what happens when you put together the efficiency gains with the Edgeworth-Salinger effect, and what you see here in the second column is that when you only have the efficiency gains -- remember that there's no foreclosure here -- prices went down by 2.4 percent.

And in the first column, I replicated the original regression, the first regression I show you, and you have to remember that the weighted effects for that regression was a decrease in prices of 0.9 percent. So we're talking about a huge effect of the Edgeworth-Salinger -- a huge Edgeworth-Salinger effect on prices when you include it in the analysis. You go from the 2.4 percent reduction to 0.9 percent.

Okay, let me talk about the additional things that we have done in the paper. The first thing we did was to look at bordering counties. So that's important because we want to have good controls for the counties that were exposed to either vertical integration or vertical integration and the

272

Edgeworth-Salinger effect. So we limited the analysis just to bordering counties that were differentially affected by vertical integration, and we find exactly the same.

Second, this is an industry where sales are very important. We see sales all the time. So we have -- we redid the analysis both just on regular prices and just on sales, and we find larger effects on regular prices, both for efficiency and the Edgeworth-Salinger effect, but also significant, very large effects when you look at sales prices. So we're basically getting the same results.

We can play quite a bit with alternative and more extreme versions of the fixed effects, basically triplicating the number of fixed effects or something like that, things like that, and we still find the same effects. So if there is something that is really robust coming out of this story, it is the Edgeworth-Salinger effect is incredibly robust, it does exist, and we should consider it when we're talking about vertical mergers.

There are, however, some alternative explanations for our findings. So the first obvious one is market foreclosure, and I already spend some time saying why, in this particular case, we don't

273

think that foreclosure is a concern.

The second one that was suggested by Paolo a month ago or something like that was, well, what if the bottlers are capacity-constrained, because you have a decrease in the input costs of Coca-Cola and Pepsi products, and if you are capacity-constrained, a natural reaction to that is to increase the price of Dr. Pepper to free capacity to produce more owned brands.

There are two things there, but the most important one is that is probably a very good explanation for short-run effects. The economic difference-in-difference results suggest that the effect is actually quite persistent over time. The other thing is that it seems like expanding capacity is not that expensive anyway in the long run.

Finally, another thing that could be happening here is, well, what happens if, instead of the Edgeworth-Salinger effect, what's going on here is that Dr. Pepper bottlers, in nonvertically integrated areas, actually change their frequency of sales. And in the paper we actually ruled that out, and what we show is that Dr. Pepper bottlers in vertically integrated areas actually increase a little bit the frequency of their sales. So we ruled that out, also.

274

So let me finish with this. We haven't said the other vertical integration actually is a trade-off between efficiency and foreclosure. What we say is if we have multiproduct firms, we have to take into account the Edgeworth-Salinger effect, and this is, to my knowledge, the first paper to actually put a number on the Edgeworth-Salinger effect.

What we show is that it counteracts the impact of efficiency gains to a large extent, and that's the reason why we believe it should be part of our standard toolkit when one thinks about vertical integration.

Thank you.

(Applause.)

MR. ROSENBAUM: The paper will be discussed by Andrew Sweeting.

MR. SWEETING: Okay, thank you. This is -- I'm very glad to be discussing this paper. It's a very clear paper. I think it's a very important paper from a policy perspective. Fernando did a great job of explaining what's in there, but just to kind of reiterate on the main points, right, so they're looking at this setup where they're focusing on kind of three firms, so Coca-Cola, Pepsi, and Dr. Pepper, and they have this geographic variation, okay?

275

So they are going to see Coke and Pepsi vertically integrating with some of the most important bottlers, and then this is going to have different effects geographically on Dr. Pepper depending on whether those bottlers also distribute Dr. Pepper in those particular counties.

Okay. So the results, which obviously Fernando discussed, is that they see the vertical integration is associated with a lowering of the prices for Coke and Pepsi's products. On the other hand, they're seeing that the retail prices of Dr. Pepper's products tend to go up, okay, and they're noticing that the percentage increase in price in the second point is greater than the percentage reduction in the first point.

Okay. So there's just lots and lots of things to like in the paper. So the theory presented is very simple, and I think that kind of makes it very plausible for a lot of different settings. So the theory they developed, which Fernando actually didn't say that much about, is in the kind of extreme simplest form in the sense of the wholesale prices coming from the syrup makers held fixed, then they're just going to focus on the incentives once there's vertical integration to play with the downstream

276

prices.

One reason why this is a good setting to look at is the beverages are pretty high-margin products, so we can think the small percentage changes in the margins are going to have potentially quite large effects on prices. The empirical analysis is very transparent. The magnitudes are pretty consistent across different specifications, and I think it's particularly nice that they're consistent across Coca-Cola and across Pepsi.

The authors actually draw a very clear policy conclusion. So, on average, the prices, at least when you're looking at the nonquantity-weighted form, go up after mergers, and, therefore, they say, you know, a standard thing that antitrust authorities should look at when they're looking at vertical mergers, even if there's not a risk of foreclosure, is this kind of Edgeworth-Salinger effect.

Okay, so here are kind of my main comments. So the paper is kind of very clean, and it's so easy to read because it's kind of short and to the point and you get to the results kind of super, super quick. On the other hand, there's -- I think, you know, the paper would benefit and the reader would benefit from having kind of more discussion of the context, okay?

277

So in this industry, what we know is that there's a history of bottlers kind of integrating with upstream firms and de-integrating with upstream firms and legal battles involving people who are bottling for other syrup makers, and before this vertical integration took place, one relevant thing is that Coke and Pepsi owned substantial proportions of these bottlers that they ended up integrating with.

Okay. So at least one interpretation of this is that the Edgeworth-Salinger effects that are going to be identified are probably going to be underestimates of the true incentives, because these incentives should already have been at play before the vertical integration that they look at.

A second relevant factor which Fernando mentioned, the fact that at the time of the vertical integration, Coke and Pepsi signed new bottling license agreements with Dr. Pepper for these distribution areas. I think there needs to be a little bit maybe more discussion about what these kind of long-term contracts that Dr. Pepper signed, how that affects how we should think about the model, right?

So the way the model and the work is currently presented, you would kind of get the impression that

278

these price changes are basically inflating a lot of harm on Dr. Pepper, whereas obviously Dr. Pepper was willing to sign these agreements. And one thing I went -- I was having a look at Dr. Pepper's earning calls around the time that the agreements were signed, and there they -- you know, and maybe unsurprisingly -- they were portraying the loss of these agreements as being very good for Dr. Pepper.

They talked about kind of performance targets that were in these contracts for Coke and Pepsi, and, in particular, what -- they referred to something -- which I wasn't quite sure how to interpret -- which was the repatriation of Dr. Pepper volume from the bottlers to Dr. Pepper.

Okay, I'm not quite sure how to interpret that, but what it makes me think is these contracts obviously are connected with partly what Dr. Pepper saw its future strategy for the next 20 years as being, and also just the length of the contracts and the very large lump sum transfers of hundreds of millions of dollars that went on probably makes you think that these incentives within these contracts wouldn't simply involve, you know, pure linear pricing, even if there was some margins on the upstream being charged.

279

Okay. I think it would be good to think more specifically about also what we see going on in areas. You know, in the control group here are both areas where Dr. Pepper is vertically integrated and areas where Dr. Pepper is distributing products through bottlers who are not owned by Coke and Pepsi, and I think it may be interesting to separate out those different areas to maybe understand, you know, were there some things that Dr. Pepper was implementing at the same time that maybe went through its own bottlers but not through independent bottlers.

You know, a lot of branding and promotion here is going to be national, so that even if it isn't -- even when there's this separation across counties in the vertical structure, it may be the case that some things that happen in the treated counties are going to be playing over to effects we see in the control group.

Obviously, one thing we observe here is retail prices, right? So the model Fernando put up on the board was vertical integration between manufacturers and retailers. Here we have vertical integration really between manufacturers, bottlers, who were then selling on to retailers, who then sell on to final consumers, and what we observe is retail prices, but

280

retail is kind of excluded from the picture here.

Now, I think probably the justification for doing this is that carbonated beverages are kind of a classic example of direct-to-store delivered products, where the bottlers in this case would maintain a lot of control over how stuff's presented in the store, you know, what goes on different shelves, and so on. But I think at least in terms of considering how we want to think about correlations and possible residuals across counties, when we have the same retailers operating in multiple counties across these borders, I think is relevant.

Okay. So Fernando already mentioned this, so I would like to see kind of more focus on quantities, right? So if we're just looking at price changes, obviously different products are sold in different quantities, and for the same product over time, more is going to be sold on sale than when it's not on sale. So I do really think we could learn -- you know, it's interesting to look at what happens to quantities if we want to start thinking about welfare.

It's also -- if you look at actually what happens in the quantity predictions, if you look at kind of the mean residual for Dr. Pepper products compared with Coke and Pepsi products, it's actually

281

the case that the Dr. Pepper products are gaining in quantity relative to the Coke and Pepsi products, and that's obviously something that's quite different from what you would see from the price regressions.

As Fernando mentioned, it is also the case that in these vertically distributed areas, Dr. Pepper actually goes on sale more often after the vertical integration than before the vertical integration. And here you can actually see that, at least in terms of national volume-related market share, Dr. Pepper's market share is not going down after these agreements take place.

Okay. So I'd also kind of push maybe on examining the distribution of prices more rarely. So when you're using the IRI or the Nielsen data, you know, it's very easy to get kind of 37 million observations. I'd just be kind of knocked dead by that, and, you know, just getting -- kind of computing the fixed effects regression is going to take you a lot of time.

But on the other hand, I think it's also important to think about, you know, what are actually the prices being charged in the store and how they may differ from the kind of average revenue measure that you tend to get in these scanner data sets, right?

282

So here I was partly thinking about this because -- so the University of Maryland is a Pepsi-only campus, but they -- you know, this is a kind of example of state-assisted foreclosure in this case, but one thing they do sell is Dr. Pepper, okay? And you might have thought that because of the absence of Coke, it would actually make the incentives to engage in Edgeworth-Salinger pricing kind of particularly strong, but at least everywhere that I've seen on campus, Coke -- Pepsi and Dr. Pepper are sold at exactly the same price.

Similarly, when I was wandering around grocery stores this weekend trying to think about how Dr. Pepper and Pepsi are actually priced, wherever I went, whether things were on sale or were not on sale, Pepsi and Dr. Pepper were being charged at exactly the same price in Montgomery County, which is one of, I believe, these vertically integrated counties.

Okay. So, finally, obviously, you know, this is a very kind of reduced-form paper in kind of a good sense, partly because that's buying us a lot of transparency. I think a structural exercise could add insights here, and really that comes in two kinds.

So you could write another paper which kind of used a structural model to really start thinking hard

283

about the welfare effects, but I think also just maybe taking the theory kind of more seriously in terms of what a model would imply about which products of Pepsi and Dr. Pepper are particularly close substitutes. In terms of what are not close substitutes to products being sold, for example, by Coke, where Pepsi and Dr. Pepper are the vertically integrated pair, I think might shed a light, right?

Are we seeing the price increases on the right kinds of products? Given the particular distribution of tastes for those products and a particular kind of vertical integration we're seeing, I think would provide nice confirmation that the story -- of the story that's going on.

Okay. So, in summary, I think this is a really good paper. I think it should have implications for policy. I think the authors have lots of scope to probe, using this data, kind of deeper into these issues, which have received very little previous attention, but are clearly very important.

MR. ROSENBAUM: Thanks, Andrew.

We have time for one or two questions.

MALE AUDIENCE MEMBER: Just a question around, so you showed the kind of effects if Coke and Pepsi were the ones who integrated with the bottlers, so we

284

see an overall back-of-the-envelope price decrease of 1 percent, but I'm thinking that if Dr. Pepper was the one who was integrating, and do you have some kind of, like, counterfactuals or some kind of thoughts that you have put on this?

MR. LUCO: So it definitely depends -- so the outcome will depend on the different market shares, for sure. We don't have anything on that. So that re -- that's what Andrew is saying, if you push these in the structural direction, we can actually go and do that kind of counterfactual, but we haven't done that.

MALE AUDIENCE MEMBER: Thanks. That was great. I was just wondering, I'm having a hard time differentiating benefits from vertical integration from benefits from anything else that might increase the bottlers' profits, like improving the delivery system from Coke, Coke's delivery system of whatever it is they deliver to the bottlers, compared to Dr. Pepper. Wouldn't that have the same effects, and, therefore, should we look askance at anything that reduces costs, not just double marginalization?

MR. LUCO: Okay, let me see if I understood the question right. You would get exactly the same results if we just talk about a retailer that faces a decrease in the cost of one of the products itself.

That's absolutely true. In this particular case, it is caused by vertical integration, so that's why we're pushing in that direction. I don't know if that answers your question.

MALE AUDIENCE MEMBER: Well, I'm just sort of wondering -- if we think we should be incorporating this effect into the analysis of vertical integration, I'm just wondering whether if some -- if Dr. Pepper came to the FTC and complained that, hey, Coca-Cola came up with a better way of distributing stuff, and because of that, the retailers or the bottlers no longer want to carry my product anymore, should we say, well, gee, that consumer welfare has gone down or those losses may outweigh the gains from the savings of the costs, and, therefore, we should block these cost reductions.

MR. LUCO: It's a tricky question. Let me put it in this way. Again, any changes in relative costs are going to cause these type of results. Whether these are -- whether technological changes or antitrust concerns, I would say the answer is no. In this particular case it's because vertical integration is causing the change in pricing incentives is what I would be worried. Yeah, that would be it.

(Applause.)

MR. ROSENBAUM: Thanks again to Igal for putting that session together and thanks to all the presenters and discussants.

We will take a 20-minute break, and then we will come back for Igal's keynote address.

(A brief recess was taken.)

KEYNOTE ADDRESS

MR. WILSON: All right. Thanks, everybody. I think we're starting to run a tad long, so I'm going to move towards introducing our final panel or our final session of the evening. This will be the -- our second keynote of the conference. So Dr. Igal Hendel will be talking to us about health insurance market design. Igal is the Ida C. Cook Professor in the Department of Economics at Northwestern University.

His research interests are in applied micro and industrial organization. Some of his recent work has touched on markets with asymmetric information and involves the estimation of dynamic consumer behavior. In addition, he has served in an editorial capacity on the board of editors at the AER and previously was a co-editor both at the RAND Journal of Economics and an Associate Editor of the JIE. Thanks very much.

(Applause.)

MR. HENDEL: Thanks, delighted to be here. Thank you for having it. It's a great conference. I really enjoyed all the papers so far.

So what I'm going to do is -- I'm going to promise you that it's going to be helpful, you know, I agree that this is policy-relevant. It's not really antitrust-related, but, you know, you are going to

tolerate it.

It's going to be mostly going over what I've been doing in the last couple of years, I'm going to hopefully be doing in the next couple of years, and it has to do with the design of insurance marketplaces, exchanges, right? So they are very -- you know, they are in the news lately every couple of -- every couple of weeks, they come back again, and so what do we mean by exchanges?

You probably know, you don't need much explanation, but they have been designed in many places, Switzerland, Netherlands, and so on, and what it means is some kind of rules for opening a market. Typically they involve annual contracts, free entry, some pricing restrictions, some minimum coverage, like we saw two papers back, and a well-defined product, you know, or products, you know, 60, 70, 80 percent actuarial value, so the customers know what they are getting -- subject to some, you know, tricks played by companies -- and that way they can compare prices that they find in the marketplace.

So what are we going to be -- so what we looked at in the past was at pricing restrictions, at prices, how do they affect participation, adverse selection, and so on. As you know, again, if you -- you know, if you watch some TV or, you know, you look at the news online, there is -- all the time there is replacing and -- you know, with alternative plans, better way -- what is it, the -- empowerment and employment and accessibility and whatever, by a bunch of Republican Senators.

And what all these proposals have in common is that they repealed the participation mandate, and so it's perceived as infringing freedom or -- I don't know what -- you know, whatever it is perceived, they want to get rid of the mandate, and some of the proposals, you know, the -- they propose an alternative, you know, participation mechanism that I'm going to try to evaluate in a moment, and some of the proposals also get away with the preexisting conditions and the pricing of those conditions.

So basically what we are doing in the project I am going to describe is sort of play with these rules and simulate the market when it changes rules to try to see how they impact allocation and the coverage in the market.

So what are the main economics behind the design of these contracts? Well, it's two types of risks, you know, that were, you know, discussed earlier today. One is the type itself, right? So you

may get a condition and you may need insurance for that -- you know, those conditions. The other one is that that type is changing over time, so one, let's call it, a reclassification risk, that over time your type is changing, and you would like insurance against that.

The other risk is conditional on whatever your type is. You want insurance for the distribution of health expenses conditioned on your type. So that generates two issues.

One is reclassification of risk if the rules of the exchange are such that health conditions can be priced, right? So if health conditions can be priced, there is no pooling, everybody gets their own individualized price. In theory, there is going to be 100 percent participation, 100 percent trade, right? There is no adverse selection because you have an individual price for you. But if that happens, it means as you age, you are going to be facing random premiums, and that is, you know, welfare-compromising.

Now, on the other hand, if you prevent discrimination, you're going to, you know, reduce or eliminate reclassification risk. Now, if your condition cannot be priced, great, you are insured against that risk, but on the other hand, if I'm

priced by the average and I'm in better health than the average person, well, I may opt out of the market and may buy suboptimum insurance, be underinsured. So that's -- that potentially could generate adverse selection. So these are kind of the main two forces affected by the pricing rules in the market. So there is a tension, right? So the more you lower pricing, the more reckless we get you on risk and the less adverse selection in the market.

The ACA, Obamacare, went to an extreme of fully banning the pricing of health conditions, so fully eliminating reclassification of risk, at the potential cost of generating adverse selection, and, you know, we do see some, or at least in the numbers from the Massachusetts Exchange from before, that the lower coverage were the most popular insurance plans.

So one question that one may want to ask is, well, to what extent should health conditions be priced? So we trace them kind of frontier, that if you fully ban them, you become reclassification risk, and you may induce adverse selection. If you fully allow them, you (indiscernible) and you induce adverse selection.

Now, how do we answer or how did we answer that question? Well, we want to compute welfare, and the

answer is going to depend on generating an equilibrium from some population -- I am going to describe in a second what did we use for that population -- on which we want preference, preference to our risk. We want to know a distribution of types. So think about this being the market, I would like for each you to know your type of health type. And I also would like to know the distribution of health expenses that you face given your type and how those expenses change over time, so that these are basically the main ingredients.

When I have all that, I can compute the amount for each person in the market. I can, you know, generate some premium, personally breaking even, see who joins, see if those who are losing money, making money, and so on, until we convert to some notion of equilibrium.

Once we have that equilibrium prediction, we can compute, you know, how much surplus is generated in the market. And, again, what would be the exercise? The exercise would be we try different pricing conditions to see where in that frontier we do an adverse selection and reclassification of risk would you maximize welfare. So that's what I'm going to show you in a second, what we found.

That previous question we answered in the context of studying contracts, one-period contracts like -- like in the ACA, but we can try, like in other places, like Chile or Germany, to see what would be the welfare consequences of long-term contracts. So both now that insurance companies are committing o insure you for one period, suppose that they sell your contract since you are, you know, 32, until you are 65, when you go into Medicare, and the idea is that that policy could, in principle, guarantee reclassification of risk from that period onwards, but at the same time, if the insurance company could price your observables, could overcome adverse selection.

So the question we want to answer is, can we get outside that frontier that I told you earlier, between reclassification risk and adverse selection, by using long-term contracts as opposed to one-period contracts. And the answer to evaluating welfare under long-term contracts is going to depend, again, on preferences of this population, the distribution of their health type, and how they transition over time.

I'm going to say if we have those ingredients -- and I am going to tell you in a second where we get those ingredients -- we can simulate optimal contracts, and we can compute welfare. So

that's basically what I am going to tell you later on.

And, finally, repeal and replace. So as you may remember, a couple of months ago, a proposal of, you know, repealing or replacing -- I don't know what it was -- at the House of Representatives entailed removing the mandate and instead relying on participation by 30 percent penalty of premiums when somebody didn't have continuous coverage, right?

So it's not a penalty for being outside, right? So there is no infringement on freedom. If you want to be outside the market, be outside the market, but if you change your mind, you're going to be penalized by a 30 percent extra premium for coming back.

So the Senate Bill had a different inducement mechanism. It was, again, full freedom. You don't want to participate, don't participate, but if you decide to come back because you have got a condition, it is going to be six months of waiting period to be covered. Both alternatives -- so, and that one, if you remember, it was McCain who voted it down with, you know, one finger, so it didn't go forward.

So both alternatives to enhance participation, it create dynamics, right, because now we have a state variable. So your choices today depend on what you did last period. So it's not that easy. It's not

that easy as solving a simple static equilibrium.

So although contracts are going to be yearly, the choice of the consumer today affects their state in the future, so we have to solve for a dynamic problem to predict demand, which together with cost is going to general that equilibrium in that market.

So the policy question here, in the context of this one-year contract, but with consumer dynamics, is going to be, well, which penalties are better? So how do you want to induce participation if you get rid of the mandate?

And to answer the question -- you can guess by now, because I'm repeating myself all the time -- what we know is preferences, we need total risk, we need the transitions across health type, and a distribution of health types, and that's what I want to tell you in five minutes or maybe ten minutes, how to, you know, get those ingredients from data that companies have been more willing lately to share, and once you have that, we can simulate other either different pricing groups in a static exchange or one-period contracts that generate demand dynamics or fully dynamic contracts in the exchanges.

And, again, I am not going to repeat myself. You know it by now. I can ask you by the end of the

talk what do you need, what are the ingredients, and I am sure you are going to know. So these are the ingredients.

So what did we use? So what we had is data from a large company. Most of you have seen, you know, prior presentations, so you already know the data, but let me just highlight what I think is interesting about that data, and what's interesting is -- and it's not unique to our data, so, again, other people have used, like, core Microsoft data.

And the key is that the data contains for each person in that population their diagnostics for at least a year. Here, it was a little bit more. If an employee stayed longer, we see the trends of how their health evolved over time because we know their claims data. We know their ICD-9 codes. So we know really what they were treated for.

Now, knowing what they were treated for and using a software, a professional software developed at Johns Hopkins Medical School, now we can forecast what this -- their actuarial value for the following year. So that is the key. So think about this, you know, for -- this is the market. For each one of you, I know your prior year diagnostics. I pass your data through the software, so now I have a number that says the actuarial cost for insuring Steve is $1,200, and for each one of you.

So now I have a whole population where I know as much as an insurer would know about that population, right? So the insurer would ask you questions, would look at your records, and would assess your actuarial type. If we have this data, that's what we know.

Now, knowing your type and knowing the ex post distribution of health expenses across everybody that looks exactly like you, now we can estimate a distribution of health expenses for somebody whose expected cost is $1,200, okay? So that way we are going to know not only your type, but we are going to know the distribution of expenses for the following year.

How is that going to help us? Well, once we have that, we can compute an expected utility. So give me risk preferences, a CARA parameter, give me a distribution of future expenses given your type, so we can compute an expected utility given any possible insurance contract. So once that insurance contract is no insurance, right, so suppose you're not insured. We know your utility, we know the distribution of costs, compute expected utility if you are going to

face a full risk.

Now, we can do the same thing if you have, like, an 80 percent actuarial value policy. That's going to deliver a different expected utility. Now, we can also compute the gap, and the gap between, say, the 60 and the zero actuarial value policies is going to be your willingness to pay for an insurance policy of 60 percent actuarial value. So it means that in this population, that I know your type and I know your distribution of health expenses and a CARA parameter, now I know your willingness to pay for different contracts that we may want to design.

If I know your willingness to pay, it means that I know the demand in this market, so I'm ready to -- sort of to simulate how people are going to behave at different premiums. Now, given that I know your type and I know how you are going to behave, we can compute the actuarial costs of offering each possible plan. So we have everything, right? So we know who's going to buy, how costly they are going to be for an insurer. We can see if they are going to break even or not, to compute some kind of predicted outcome for the market. So that's basically what we are going to do, and there are different assumptions.

So this is just summarizing what I tried to say

without an accent. Okay, data. To prove existence, there is data. These are -- again, this is ages/states. We are going to partition the states from healthiest to sickest, just to have enough observations in each cell, and as you see, the population unfortunately is getting sicker as they age.

We are going to have transitions for each age group. We are going to have how they transition from being healthy to less healthy and then, you know, luckily some people will go back. So we are going to use that when we compute the expected utility from a long-term contract, right? It's going to depend on how you transition into the future, right? So persistence is going to be key to compute welfare.

Okay, this one is neater. So here what we have is a 30-year-old in each of the possible health states from one, healthiest, to sicker, and their expected -- as we roll forward, this mark of probabilities of transitioning, what is their expected health expenditure? So what you see is that early on there's a lot of information, it sort of evaporates after five, you know, seven, or ten years, and everybody looks very similar.

Now, we find that encouraging because typically

life insurance underwriting in both information of stuff that happened in the last seven to ten years, so it means that the actuaries think that this information evaporates unless it's really a chronic condition.

Okay, so what else do we need? We need a, you know, solution concept. We are going to think of breaking even premiums, Riley equilibrium. For the contracts, I am going to do some dynamic contracts in a competitive industry when I show in a sec.

So part one, one-period contracts, pricing rule, what did we do? This is, you know, an old paper, 2015, and so this is what -- we did play with different rules allowing for more and more price discrimination. That eliminates adverse selection but induces more reclassification risk.

What did we find? Well, because of adverse selection was of the order of $600, so if you fully forbid pricing health conditions, it's going to compromise welfare. Around how much? Well, $600, which was around 10 percent of the actuarial cost -- of the average actuarial cost from the sample. So it is substantial, but it was nothing compared to the welfare loss when you start allowing for more discrimination.

When you start allowing for more discrimination, now people face reclassification risk. These are bigger stakes, so now if you are in bad health, if you are in stage seven, your premiums are going to be $18,000 as opposed to $1,000. Those are losses from welfare were way, way, way larger.

So our conclusion, our take-away was the ACA did well in banning pricing of health conditions because what they overcome is relatively small. The distortion for adverse selection is small relative to what they would -- the welfare loss from reclassification risk.

Long-term contracts, so what are we doing -- so this is current work, and so now what we want to consider is instead of one-year contracts, assume a competitive industry that offers to insure the patient, you know, for the rest of their life until they transition into Medicare.

Now, this is going to be a problem if there was two-sided commitment, right? So with two-sided commitment, we just sit together when I'm 25, 32, we are going to get the -- sort of the efficient outcome, and there's nothing to solve, and I wouldn't bore you with that. So what we think is relevant is not that two-sided commitment. Probably what we think is

relevant is a one-sided commitment problem where the company guarantees that they are going to insure you in the future, but the customer can drop the moment they want, all right?

So here I have the agencies who tell me if otherwise the contract is going to be legal or not, but I understand that phone companies, cell phone companies have trouble imposing fees from terminating coverage. So they look at me, like, what does this guy -- anyway, so here this is going to be one-sided commitment. I don't know if it's -- I think some California courts found the fees -- the termination fees illegal, basically because they -- the customer is imposing no damage on the company, so how do you justify that you -- you just tie them forever? So I don't know.

Whatever it is, let me justify on practical reasons, every insurance -- life insurance contract that we are aware of in the U.S. and Canada is under unilateral commitment. There are no penalties for dropping coverage. So that's going to be my assumption, given that nobody complained? Good, nobody complained, so I am going to -- that's going to be my assumption from now on.

Now, on the unilateral, one-sided commitment,

the issue is going to be that although the company can offer you insurance forever, it's customers in good health that are going to drop coverage, and that is going to make insurance unravel. So the optimum contract or -- the competitive equilibrium is not going to involve full insurance against reclassification risk.

So basically this is -- let me skip notations. So here the only thing I want to say is we solved it for 40 periods, from age 25, post college, until Medicare. From the data, we have these Lambdas, your type, your health type. From the data, we know what's the distribution of health costs given your type.

We assume symmetric learning. As you age, if you want coverage, you have to show up to an insurance company, and they can look at your records. They ask you to fill out questionnaire. So I'm assuming -- we are assuming that this is symmetric learning.

And what we do is we solve for the competitive equilibrium. And, again, the key are these transitions. That's how your health is going to move over time. Now, if these transitions were kind of completely persistent, once you're 25, you're either sick and you are going to remain sick or you're healthy and going to remain healthy, then the

long-term contract doesn't do anything, right, because it means information is already revealed. If you're already sick, right, so there's nothing to insure you for the future, so it's good that there are transitions across states over time and that the information is not fully persistent, as I show in the previous picture.

So what do we find? We find that optimal contracts offer a minimum consumption guarantee, so if you want to think of -- you know, around it, they offer a premium that basically is not going to go up. So if you -- if you develop a condition, they guarantee not to price you against that, okay?

But instead, if you remain in good health, they are going to give you a break. So think about, you know, for those that are chairmen here in your departments, so that's exactly what happens with -- when people in your faculty publish well, right? So if they didn't publish, they are stuck and you have to keep paying them the same amount, but if they publish well, they have an outside option, an outside offer, and you have to match that higher outside offer. That's exactly what's going on here.

So here what's going on is if the person is in bad health, the long-term contract insures them forever. The premium is never going to go up. Instead, if they are in good health, they can approach a competitor that's going to give them a better premium to reflect that they remain in good health, and with that premium, they can go back to the original firm and say, look, I have a better premium, lower my premiums, and that's exactly what the optimal contract does.

And, you know, this is the counterpart of Harris and Holmstrom in the labor context, sort of in the chairman context. Somebody who proves to be more productive gets a better deal. Somebody who proves to be unproductive does not suffer a wage loss. So basically that's the nature of the contracts.

The consequence is that this optimal coverage cannot fully guarantee against reclassification risk. The 25-year-old knows that they can insure against bad drugs, but they cannot lock themselves in into the policy if they happen to be lucky and healthy. So for that reason, they cannot equate marginal utilities across all the future periods on states, right, because if they are in good shape, they are going to have better deal, and they cannot transfer resources from that good state to a bad one, okay? So there's partial insurance against reclassification risk.

So what do we do? We simulate the equilibrium with our CARA parameters, some discount factor, under competitive assumption, the seven health states. What do we get? Let me be very brief, because I'm sure you want to run away from here soon. For a flat income profile -- so the optimal contract depends on the profile of income.

For a flat income profile, so somebody whose wages increase at the same pace that the medical costs of that -- you know, of that group increases, so basically the net income is flat, it means that is a population without any saving and borrowing. I want to neutralize that so we don't -- so the insurance company doesn't become a bank.

So what we have there is on the left, is the first pass. The first pass would be like 53.67 thousand dollars, and that will be sort of the welfare -- the monetary -- the money metric of the welfare of a person that manages to consume their first base allocation, right? They get full insurance against their medical and against their type risks, okay?

Now, the second number, 52.47, is the welfare that same person would get if there are just one-period contracts that are fully priced. That's,

if you want, the (indiscernible) of the spot market. So you open the -- this is before Obamacare, there were just spot contracts, and there were no -- it's not true, but suppose that ban on pricing risk systemization is removed, insurance companies can price whatever they observe, they are going to get full trade everywhere that gets insurance, but their premiums are going to be jumping around over time.

So that would be the second number, 52, and you see there is a loss of $1,200 associated with the risk that that premium reflects, all right? So you have full insurance against a medical risk, but your premium is jumping around. So the welfare loss of $1,200 comes from that risk, from the reclassification of risk, comparing these two numbers.

Now, the third number is certainty equivalent D, for dynamic, that would be the certainty equivalent if companies are able to offer dynamic contracts, and what you see is that it goes quite up, almost all the way to first pass. So it appears that dynamic contracts are great, but I'm tricking you.

So the reason I'm cheating is that I'm computing that where somebody was flat net income, right, so with somebody who has enough income early on in life that they're willing to put money in that

contract to subsidize their future selves, and obviously we -- nobody, right, is that miserable in the market to get the same wage at 60 as when they were 25. So in a way I'm using it for a population that doesn't reflect many workers.

If you see the bottom -- the one at the bottom, in a second, that's what we call a manager. It's a person in our sample who has a much steeper income profit, that's a real income profit, and you see that the gain from -- there is a gain from long-term contracts, but it's much -- it like goes how long -- how much, like a third or two-thirds of the way. So it's not that effective.

Why? Because this person is poor when he's young, so he's not willing to put that much money up front to pay for the future premiums. So dynamic contracts help, but it depends on the income profile of the worker.

The final number is the ACA. Now, what is the ACA? Well, the ACA is open the market with static contracts, and the reason that number, 52.87, is under the first list is because of adverse selection. So if you compare the first column to the last one, that is the loss from adverse selection.

Now, in this particular example, dynamic

309

contracts do a little bit better than the ACA for the manager and do a lot better than the ACA for the flat income, but remember, the flat net income is a fictitious worker. So we should look at the second one, and the second one is a very minimal gain from dynamic contracts, not that much better than the ACA.

Finally, the Republican reform. What do I want to say about the reform? The reform, we go back -- this is, again, ongoing work, so Mike would be very upset if he knows I am even mentioning this, because the numbers I'm going to show you at the end are fake or are, you know, very preliminary numbers that -- you know, don't tell him. Anyway, I am going to share the numbers, but don't tell him.

So what we're doing here is going back to one-period contracts, that because the Republican proposal involved future consequences for today's actions, now we have a dynamic problem, and what we do is we solve that dynamic problem -- and let me skip notation -- but basically this is sort of a Dixit model, that you are either out or in. If you are out and you want to have -- want to go back in, you have to pay a fixed, you know, penalty, and so on.

Solving that for a vector of premiums from age 25 to 64, we get the value functions. Once we get the

310

value functions, we know the month. Once we know the month, we can compute costs, and keep iterating until we find a breaking even vector of premiums, which in principle is going to be in equilibria.

Once we do that, we get these numbers that I just -- I just want to show they exist, but now I'm going to hide them. Sorry. Okay, I am going to show you, but do not forget. The only thing I want to highlight there is the House proposal with a 30 percent penalty, that's very similar. Again, we saw our -- anyway, I shouldn't -- I shouldn't show you this.

This is a very first cut that we are not very proud. It was just kind of the first, you know, simulation we did just to entertain you. Do you see the House 30 percent, that's very similar to the 4 -- to sort of the ACA kind of $400 penalty, roughly, that if not -- so, but on the other hand, the Senate proposal that keeps you -- and this is -- so we couldn't wait half a year, so we waited a whole year, so I keep, you know, cheating here, but anyway, what you see there is participation is almost 100 percent because people are really, because they are risk averse and they are so in panic of developing a condition and not having coverage, that most of them

311

participate.

Now, what you see -- I find it interesting -- is that for the, you know, older people, that loses value, right, because the horizon gets smaller and smaller. So if you think a 64-year-old has an option value, so those people, you know, sort of pull out of the exchange, but for younger, they are paying for the option value of remaining. So, again, I didn't show you the numbers, and let's move on.

But basically that was kind of the take-away, that if you believe that mandate is an infringement on liberty, there are ways to induce a participation, and it really depends on the details. So, you know, the policy, I think they are important to create sufficient participation.

So what did I say? I tried to say that there is plenty to be simulated, treating health insurance policies as financial instruments. The nonfinancial instrument could be accommodated, but in our framework, we don't have data on that. Using data that is becoming increasingly available, hopefully, you know, the Government could, you know, help us, you know, get more data.

And, again, this is the magic. This is what we are most proud of. Again, it's not my -- it's not we

312

are first to use it. I think Bob Town used it first and then Ben Hallo (phonetic). This helps you have as much information as an insurance company would have on the market.

I think food is ready. Thanks.

(Applause.)

MR. WILSON: Thanks very much. I do think we have time for just a question or two before we retire to drinks and snacks.

Anybody?

MR. GERUSO: So this idea -- so I really liked the whole research agenda with the long-term risk, reclassification, and thinking through those things, but in terms of long-term contracts versus mandate and subsidy, I mean, I guess your experiment is imagine we can't have a mandate because it's politically unpalatable, how does long-term contracts do?

But here it seems like behavioral -- so nothing in my papers has ever acknowledged behavioral economics exists, so -- but, I mean, in thinking about, you know, long-term contracts versus mandate and subsidy seems like the behavioral factors there could be pretty important, right? So someone believes that they're invincible when they're 20, just -- I mean, you -- like, actually, like, in implementing either a long-term contract or a stay out of the market for a year thing, I mean, we might -- if we implemented such a policy, we might very quickly decide we don't like it, and a couple years later, we would have a new ACA because people are not making these forward-looking decisions, and, therefore, while they're out of the market for a year, a cancer is metastasizing in them. So just -- do you have -- can you say anything at all about that?

MR. HENDEL: Maybe. So I can tell you we have bigger problems than that, if that's an answer. So currently a market like that would be co-existing with the employer provider, which dominates for tax reasons. So I think for practical purposes, an individual wouldn't like to frontload and then a year and a half later to find employment, much (indiscernible) they want to.

So the only excuse I have is that there were products like that attempted on the market, where the insurance company offers you an option of coming back, especially if you can prove the reason you're dropping is because you found employment, you don't lose your savings, if you will, you don't lose what you frontloaded. You will be taken, say, later.

Now, again, I am 100 percent with you that for

behavioral reasons or whatever reasons, many people are declining even free insurance at the moment. So honestly, just any rational model is not going to capture what's going on, and they don't know what's going on. So that's all I can say.

(Applause.)

MR. WILSON: Thanks very much for an interesting day. There are drinks and snacks back where food and coffee was earlier today.

(Whereupon, at 4:15 p.m., the proceedings were adjourned.)

CERTIFICATE OF REPORTER

I, Jennifer Metcalf Razzino, 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.

s/Jennifer M. Razzino
JENNIFER M. RAZZINO, CER

D

F

J

O

In the Matter of:

10th Annual FTC Microeconomics Conference

November 3, 2017
Day 2

Condensed Transcript with Word Index

For The Record, Inc.
(301) 870-8025 - www.ftrinc.net - (800) 921-5555

FEDERAL TRADE COMMISSION

THE TENTH ANNUAL

FEDERAL TRADE COMMISSION

MICROECONOMICS CONFERENCE

DAY 2

Friday, November 3, 2017
9:00 a.m.

Federal Trade Commission
Washington, D.C.

FEDERAL TRADE COMMISSION

I N D E X

PAPER SESSION

MR. RAVAL: All right, everybody, we're starting the first session of the day. So this is the paper session chaired by Steve Berry. So the first paper we have is An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program by Vivek Bhattacharya at Northwestern University.

MR. BHATTACHARYA: All right, well, thanks a lot for having me here. It's been a really fun conference so far, and hopefully that doesn't change with this paper. I'll be talking about an empirical model of R&D procurement contests, and I'll be using it to study data from the Department of Defense.

Okay, so, the starting point of this project is that competition plays a nontrivial role in R&D-intensive markets. If you increase competition, then you change these firms' incentives to exert effort to invest in R&D in this case, and that, in turn, can influence outcomes. And it can do so possibly adversely. So, ex ante, it's not clear that more competition necessarily leads to better price, better quality, better social surplus, better consumer surplus, or anything else we might hear about. And, of course, this nontrivial relationship between competition and innovation has led to a large

empirical literature, a large theoretical literature.

I'm going to focus in this paper -- I'm going to contribute to the literature by looking at a very particular type of R&D-intensive market. I'll be looking at what I call an R&D contest. Now, I'll be a bit more clear about what I mean by that, but this is loosely a setting where a bunch of firms are competing with each other to develop some sort of innovative product and then supply it to a procurer.

And they often compete over multiple stages. In my case, they are. They're going to do that. You can think of these stages as loosely consisting of an initial research phase where you get a sense of what you can build and how much the procurer would value it. If you're successful there, you can actually go -- move into develop and try to build what you said you'd build. And if you're successful at both those steps, you can compete with the other firms in the contest to deliver the product to the procurer.

In my case, the procurer is going to be the Government, the DOD in particular. And one nice thing about -- one thing about these sort of government procurement contracts or R&D contracts is that they're often structured in a way that looks like a contest and that there's a winnowing down of R&D contracts over time and there's -- there's -- and these contracts lead to some sort of procurement, either implicitly or, in my case, explicitly.

So the broad question I'll ask is how do they extend to competition and more generally the design of these contests affect the outcomes that we see.

And there -- so in order to do that, I'm first going to make a methodological contribution, that there's a fairly sizable theoretical literature on R&D contests, but relative to the theoretical literature and relative to the empirical importance, there hasn't been much empirical work trying to understand the sort of heterogeneity that governs the outcomes we see in these contexts.

So that's what I'll try to do. I'll write down a fairly simple model of R&D procurement contests, and I'll be very clear about what features of the data identify the primitives of the model. And throughout the paper, I'll be looking into the particular Government program. I'll be looking into the DOD Small Business Innovation Research Program.

All right, so -- and this is essentially the structure of the program. I'll walk through it step by step. So every year the DOD lets about a thousand

solicitations for fairly narrow projects. So these are like widgets for airplanes. A couple years ago, the Navy was looking for something called a compact auxiliary power system for one of their amphibious combat vehicles, so this is essentially a battery that has to satisfy a set of specifications. Any firm who wants to build that battery or first develop that battery and then build it can submit a technical proposal to the DOD, and the DOD is going to score these proposals and let a few of these firms move on to phase one.

And phase one is essentially where the contest starts. This is a quick-and-dirty phase where firms get some R&D contracts from the DOD, and they do some preliminary work to try to figure out how to make their project technically feasible. Okay?

When I take this to the model, I'll think of this as a setting where these firms are going to exert effort and get a draw of value. They're going to get a sense of what -- how many features of the battery they can actually satisfy and how much that would be value -- how the DOD would value that.

At the end of phase one, these guys write another technical report. They extend it to the DOD, and the DOD is going to select a subset of these firms

to move on to phase two. And phase two is really about develop -- reducing development costs. So this is figuring out how to actually manufacture the product or deliver the product at the minimum cost possible.

This phase tends to be much more intense, and they're -- and these -- they're larger R&D contracts, and there's a lot more variation across firms and the size of the R&D contracts they get.

In this phase, I'll think of -- in the model, I'll think of this as exerting effort to get a draw off your delivery cost, and the feature that I'm going to -- and when I write down the model, I'll take into account that these guys are receiving R&D contracts and that there's a limited number of spots in phase two.

Finally, at the end of phase two, if the DOD is satisfied with one of these projects, they can actually contract with the firm to do delivery. And, so, phase three is essentially a delivery phase. When I take this to the model, I'll think of this as some version -- as the contract price being set through some version of Nash bargaining, which effectively means that firms are going to expect to capture some portion of the surplus.

So I'll write down this model formally in a couple of slides. I'll show you how to identify the primitives and estimate them, and I'll use them to quantify the inefficiencies that are embedded in this setup. Okay, and you can already get a sense of what these inefficiencies are going to be, at least qualitatively.

There's something like a holdup problem in that firms are going to capture a portion of the surplus, not the full surplus, so that means they have less than the socially efficient incentives to exert effort, but counteracting that are something like a business-stealing effect, where if I displace someone from phase two, I capture their full profit, so that gives me more than the social incentive reason to exert effort. And there's also going to be something like a reimbursement effect in that these R&D contracts are going to be socially neutral transfers, but I'm going to treat them like prices as referred.

And understanding these inefficiencies are going to help us understand some simple design counterfactuals that I'll talk about. And in particular I'll focus on changing the number of competitors. If you add another competitor, then, sure, you get another draw for the pot, but now everyone else realizes they're facing more competition and there's an indirect incentive effect of exerting effort.

And I'm going to try to quantify the defect and see whether -- whether adding competition is actually beneficial in this setting. And I'll also talk about changing the intents in margin of competition, if you will, by changing the surplus that you commit to give these guys in procurement. I won't have time to discuss other design changes today.

So the data that I have comes from the Federal Procurement Data System, so I have all Navy SBIR contracts from 2000 to 2012. There are a number of reasons to focus on the Navy, but for now, just worry about the data reasons, that they were nicer with data for the most part. There are -- so I have the number and the identity of competitors at each stage. I have the R&D contract amount at each stage. If there's a phase three procurement contract, I see the contract amount as well.

Now, these projects are somewhat heterogeneous, so I'll try to control for that as best I can by looking -- by getting program-level characteristics -- or project-level characteristics from the Navy SBIR program office. And, so, I see the

contract duration, the fiscal year, the division of the Navy that developed the project, the acquisition program the project is a part of, and I'll also see the full text of the solicitations and the abstracts of the winning proposals. So that's about 15 or 20,000 pages of material. I'm going to run that through a fairly off-the-shelf machine-learning algorithm and essentially generate topics for each one of these, these contests, and try to control for heterogeneity at that point.

And here's some examples of topics. You can go much finer than that, and it does a pretty good job. I mean, the algorithm is built for stuff like this.

So I'm just going to give you a quick taste of the data without going into much -- any sort of detail about correlations. These tables show -- oh, this table shows you the distribution of the number of competitors at each stage in the contest. As you can see that these are fairly small contests. They're usually about two to four competitors in phase one. About 17 percent of contests don't even make it to phase two. The DOD says that they're not satisfied with the research done in phase one, so they just end the contest there.

About three-fourths of these contests are conditional on making -- making it into phase two, and about three-fourths of them just have one competitor. And the rest tend to have two competitors.

Phase three, the acquisition phase, it actually is fairly unlikely. So only about 10 percent of contests make it all the way to a delivery contract. Okay.

So I'll think of this as essentially -- I'll model this as research being some -- having some sort of stochastic component, and, in fact, failure rates are going to be important in identifying the primitives for me.

The other key observable are these contract amounts. So I'll think of this as measures of R&D expenditures. So if I were to plot a distribution of phase one contract amounts, it would be a mass pointed at $80,000. That amount is essentially institutionally set, effectively across the Federal Government, but there's actually a lot of variation in phase two and phase three contract amounts.

So the first distribution shows you that the distribution of phase two R&D contracts can be as low as $250,000, can be high as 1.5 million or so. The delivery contract can be as low as a few million

dollars and can go up to about 50 or 20 million, sometimes even a bit higher.

Okay. And that variation was due to cross-project variations, or the Navy cares more about certain projects, but it's also due to variation in who shows up to a project. So that's what that final histogram shows you, that's the percent difference between the -- I'm looking at contests where multiple people show up to phase two, so I'm perfectly controlling for project-level heterogeneity, and I'm looking at the percent difference between the highest funded guy and the lowest funded guy. And that number is often between 25 and 50 percent, sometimes larger than that.

I'll interpret that variation as variation in value. That's consistent with a bunch of things that I discuss in the paper. It's consistent with what the Navy discusses or attributes that variation to. They say that they give more funding to projects with higher transition potential. It's consistent with descriptive correlations of project-level success.

Projects with higher funding tend to succeed at higher rates, tend to lead to delivery contracts at higher rates, both across projects controlling for a

13

bunch of stuff within project, perfectly controlling for heterogeneity. They lead to higher phase three funding amounts. So there are many reasons to think that this is indicative of value and that that's sort of the stance I'm going to take when I take the data to the -- the model to the data.

And this is the model that I'm going to take to the data. It's actually sort of scary having a countdown clock staring you down. This is the first time I'm presented with a countdown clock.

So this model is -- it's fairly simple. It fits on this one slide. So I'll just walk through it step by step. In phase one, they're N1 firms. I'll think of them as ex ante symmetric. They each exert some effort, B. That's a probability, that's a normalization at some monetary cost I of P dollars. Okay?

Generating an effort, P, means that they generate a success with some probability, P. If they're successful, they get a sense of how much the DOD would value their project. That's draw v from some distribution f. The DOD is going to score these projects. They're going to see the Vs, and it's going to let the top N2-bar firms move on to phase two.

And if fewer than N2-bar firms succeed, then

14

not everyone moves -- just the guys who succeed move on to phase two. Now, in phase there, there are N2 firms. They each draw -- they each have a draw of v, and they're going to exert some effort, t, to draw some delivery costs, c, from some distribution, h, that's parameterized by t. Okay, and t is in dollars. That's a normalization.

Now, N2 is public. That's consistent with how the DOD announces stuff, but firms are not going to know each other's values. They're going to have beliefs or values. These beliefs may or may not depend on their value, depending on whether or not there's selection in that particular setting.

Okay. And at the end of phase two, you have some firms. They each have a v, they each have a c, and the DOD is going to see the surplus that each firm would generate if they were to bring -- deliver the product, and it's going to go to the firm with the highest surplus and pay a cost plus contract. It's going to cover the firm's costs and pay them a fraction of the incremental surplus he generates. And he's going to do that as long as v is larger than c. Okay, so there is some sort of selection condition embedded into the model. And, so, essentially phase three is something like Nash bargaining. That's how

15

I'm modeling the procurement stage here.

Okay, so you can solve for the equilibrium here. I'm looking for symmetrical equilibrium. It's characterized by a set of integral equations that are fairly easy to understand, but the important part of this slide is that the -- I'm going to make an empirical assumption that the R&D contract that I see in phase two corresponds to this firm optimal amount. Okay, it corresponds to this equilibrium.

And that's a bit of a strong assumption. That's saying that the DOD decides your R&D contract is based on what's optimal for the firm. You can justify this in a number of ways. Maybe if the DOD were to give the firm more than the optimal amount, then the firm -- giving the imperfect monitoring, the firm could try to reallocate some resources, try to pocket the rest of the money in some way. If the DOD were to give them much less than the firm optimal amount, it would be running ex post losses. That might not be great for program participation.

But I do understand that it's a bit of a strong assumption. What the important part of that assumption for most of the identification and most of the estimation is that this means that the phase two award amount is increasing in value. So this is my

16

interpretation for the DOD saying it's -- it gives higher funding to projects with greater transition potential. If you don't want me to assume this equilibrium, I'll show you I can still identify a lot of stuff about values and costs, purely from monotonicity.

And that's what I'll try to go through in a couple of minutes. Identification of this model, identifying distribution of values, distribution of costs, and the bargaining parameter, it's going to leverage three features. It's going to leverage this monotonicity thing. I see the distribution research efforts. I need the distribution values. Now I know there's a one-to-one function between them. I just don't know what that function is yet.

I'm also going to use the fact that there's a selection condition here. The DOD's only going to contract with the firm that -- with a firm that has a -- that generated a positive surplus. So if the DOD just didn't contract with the firm, I learned something about what the surplus was.

Those two assumptions are going to give me a lot of information about values and costs. In order to identify the bargaining parameter, I'm going to have to leverage the equilibrium of the model. And

I'm going to have to say that somebody's optimizing something, and here it's going to be the firm optimizing its research efforts.

And I'm going to walk through the identification proof because I think it's fairly straightforward, and it helps understand where the estimates are coming from. So in phase two, I see the phase two research effort, t, and the joint distribution of that with the phase three contract amount. Okay, I need the value distribution, f, the delivery cost distribution, h, and the bargaining parameter, Eta.

Why do I care about this? Well, the value distribution is going to tell me how much heterogeneity there is and what happens at the beginning of phase one. The cost distribution is going to tell me how much heterogeneity there is and what happens in phase two, so it's going to help me understand where competition might be useful, which phase of the contest.

So condition on a particular value of the research effort, that's like conditioning on value. I just don't know what that value is yet. I see the distribution of phase three contracts. Right now that's sort of meaningless because it's a combination

of values which I don't know and costs which I don't know and a bargaining parameter that I don't know.

Okay, but what I do know is that the contract that the DOD was just barely willing to accept is one where the delivery cost equals the value. So if I were to see a lot of these contests, the maximum value would be -- the maximum contract value would be the one where the contract amount is the value, okay, where basically they were just barely willing to trade.

So I've identified values off the support of the phase three contract distribution. This looks very stark. You can make it less stark by adding some unobserved heterogeneity. I talk about that in the paper, but this is the rough intuition, and in a stark model, this is the formal proof.

So identify values off the support, and so the value distribution is identified off the support as well. And once I have values, the residual is due to cost, so any residual variation in the contract amount, conditional on the bargaining parameter, is due to a variation that happens in phase two. That's costs.

Okay, so that helps me identify values and costs. All I've used so far is this monotonicity

argument and the selection condition that I'm only going to see a contract if values are larger than costs. I haven't used anything about anyone optimizing anything yet. But I haven't recovered this share of the surplus.

Here, I can go back to the firm's first-order condition, and note that I know everything in that equation except for the bargaining parameter. Okay, so that's one equation and one unknown, some hand-waving and some math behind the scenes shows you that there's one, one solution. And loosely what that means is that, well, where this is coming from is that from values and costs I have a sense of the marginal benefit of research, the marginal cost of doing a dollar of research is a dollar, any wedge between the marginal benefit and the marginal cost has to be due to the fact that the firm realizes they're not capturing the full surplus. Okay, and so that's what I'm interpreting as the firm's bargaining parameter.

So this is identifying a bargaining parameter off some sort of ex ante investment, which is a bit different from how at least in a conceptual sense and from how other papers that identify bargaining parameters operate, like Ali's paper or Matt's paper, but this ex ante investment is sort of a

hallmark of R&D, and I hope that that's -- this is one of the observations that could be used in other settings.

Okay, so identification hopefully was transparent. It's more robust than you think. There are a bunch of extensions in the paper, many of them -- one of which is actually relevant for estimation. And it leads to a fairly tractable estimation procedure.

Let me run through this really quickly. The loose idea is that given monotonicity I can essentially -- conditional on a bargaining parameter, I can essentially estimate the model without ever having to solve it at all. And the benefit is that that's tractable. It's not hard to solve the model, but it's not easy either, so it helps to be able to not do that during estimation, but it's also conceptually robust. So, once again, if you don't like this monotonicity, if you don't like this equilibrium assumption, you can -- you can estimate everything without actually having to impose it.

Okay. When I take -- when I actually take this to data, I'm going to have to add in some unobserved heterogeneity -- or observed and unobserved heterogeneity. Those are the covariates that I identified a couple of slides ago, the share of -- the division of the Navy and the stuff from the machine-learning algorithm. I'm also going to add in a degree of unobserved heterogeneity. That enters into a somewhat -- enters into a setting in a somewhat restrictive way.

I'll let the -- the identification proof didn't leverage any sort of cross end restriction. Different -- you might be worried the different contests -- that the Navy selects different numbers of competitors for different types of contests. I'll try to allow for that by parameterizing some of the primitives by the number of competitors in phase one, and I'm going to avoid using the $80,000 in phase one in estimation just because that's sort of an institutional number that isn't really representative of much. I'll use that as ex post check of how sane my estimates are.

And estimation essentially first proceeds by backing out the distribution of the unobserved heterogeneity in a way that's similar to Elena's paper, and then I'll do the MLE procedure that I scanned through in the previous slide, and then after I've estimated values and costs, I'm going to actually solve the model at that point and then estimate the

bargaining parameter by imposing the structure of the model at the very last stage.

And, so, here's what we learn from this procedure. I'm showing you the distribution of values as a function of N1. The Navy value seems to value these things at around $11 to 50 million, but what's more interesting is that if you take a guy from two-point-fifth percentile and you move him to the 96-7-point-fifth percentile in values, you only increase his value by about $1 or 2 million. That's around 10 to 15 percent of the mean.

So I'm estimating a fairly narrow distribution of values in phase one. Values are essentially -- and this is sort of consistent with the idea that the Navy has spelled out these projects pretty well already. Where is this coming from? The idea -- this is essentially a soft upper bound of the phase three contract distribution as a function of phase two research efforts. There are a number of contracts that had low phase two research efforts that ended up having fairly high phase three research -- or phase three procurement contracts. So that must have meant that these had high values when you -- through the lens of the model.

Costs tend to be about $7 million

conditional on it being a reasonable cost, but there's a lot more variation in the distribution of costs here. So there -- a lot of the uncertainty and research happens in the second stage. This comes from the residual variation in phase three contracts conditioning on the phase one value distribution -- or research effort distribution.

And the DOD seems to be providing these guys with fairly high-powered research incentives. Firms are acting as if they capture a good share of the surplus. All right, and if you're interested, the implied phase one research cost is about $30,000, which is not $70,000, but it's in the right ballpark, and some unobserved heterogeneity might make $70,000 somewhat reasonable as well.

Okay, so those are the estimates. With these estimates in mind, we can sort of figure out whether R&D efforts are less or more than socially optimal. In phase one, there are multiple effects at play that I discussed at the beginning of the presentation. It turns out that phase one R&D is excessive in the setting in equilibrium. The social planner would want these guys to reduce their efforts.

It's a fairly small effect when there's no business stealing. If there's only one guy, there's

no one to steal the business from. The gain from moving to the efficient level of effort is only about 4 percent. When there's a lot of business stealing, though, this can be fairly large.

Phase two R&D in this model turns out to be unambiguously less than socially efficient because firms are only going to get compensated by a fraction of their marginal contribution to society. So you can show that that means that they're always going to be less than -- they're always going to exert less effort than we'd want them to. In fact, 40 to 50 percent less effort than we'd want them to, and the surplus can be improved by about 5 to 10 percent here by sort of alleviating this holdup problem.

Okay, so what does that mean for counterfactuals? So this table shows you how -- so I'm looking at a set of parameters. If you just have one guy in the contest, then social surplus is about $140,000 in expectation. And the table shows you the change in the social surplus from change in the number of competitors in phase one and the number you let into phase two.

So if you increase the number of competitors in phase one and you still let only one of them move on to phase two, then you have a number of effects.

You get more draws from the pot, but we already estimated that those draws were fairly useless because the distribution of values is very narrow. Now, you guys -- you also have more -- more duplicate research efforts. That's socially costly, but these firms are also adjusting their research efforts downwards. And that happens to be socially rather beneficial because we estimated that they're large inefficiencies in phase one in terms of excessive research effort.

Okay. On net, that means the social surplus tends to be about the same or decrease by a bit. If you increase competition in both the phase one and phase two, then you're leveraging the fact that you see a lot of variation outcomes in phase two. So ex ante, that means that there's low substitutability across products in phase two. So if you want one guy, you want another guy loosely, so the social surplus increases pretty strongly when you add competitors in phase one and phase two.

The takeaway is that the planner prefers to invite contestants to in both stages of the contest. And the main benefits are from the direct effect of more draws in phase two and the incentive effect of these guys adjusting their research efforts in phase one.

Okay, how does social surplus -- what does social surplus look like? As a function of the share of the surplus you give the firms. So where is -- this is varying how much you give the firms with the point to the right being giving everything to the firms, and the dotted line that you may or may not be able to see is where we are. It turns out holdup costs are fairly low here. We estimated that a couple of slides ago. So there's a beneficial to -- and there's a benefit to sort of reducing the share of the surplus you give the firms and then just economizing on the other inefficiencies here.

Okay, the net benefit turns out to be fairly small. And as an aside, you might -- ex ante, you might have been worried that the DOD is giving these firms too low a share of the surplus by essentially not giving enough incentives to exert any effort. That doesn't turn out to be the case here.

Okay, so just really quickly, why don't we actually see this in practice? It turns out that many of these socially beneficial design changes are actually privately harmful for the DOD because the DOD doesn't capture a large share of the surplus. They end up paying out their R&D contracts. So they're seeing -- at least at the estimated parameters, there

seems to be a tension between what's optimal for the DOD and what's optimal for the social planner. So we might be somewhere in the middle. With more time, we could have had more of a discussion on other things that these estimates might have implied about the objectives.

But let me just end there. I did what I said I did, and I think that the takeaway is that -- that I'd like to apply to other papers is the observation that these R&D contracts are -- R&D efforts are indicative of what happened in the past and also indicative of what these firms expect in the future, so that gives you a good deal of information about the parameters of the model. And that's what I'm leveraging in this paper.

Thanks a lot.

(Applause.)

MR. RAVAL: All right, we have Elena Krasnokutskaya from Johns Hopkins to discuss this paper.

MS. KRASNOKUTSKAYA: So first of all, you know, I would like to say that this is a very interesting paper. I enjoyed reading it and at the end of the day think I've learned new stuff from this paper. Okay, so what is the paper about? So the

paper uses the data from this SBIR, I guess, program, that runs R&D contests on the topics of interest to DOD, and, in fact, on topic of interest to the Federal Government.

So this program is specifically designed to fund the research by small businesses. And the funding is allocated on a competitive basis, so at every stage of the contest, the participation is competitive -- selection into participation is competitive. And the goal here is to have a -- to have products which could be sold -- eventually sold to the military in the military market or in the private sector, right? So that's kind of the program that we have here.

So the way the author thinks about this environment or the way -- the way he studies this market, he basically writes down a model which links eventual profitability of this invention to the competitive pressure in the contest and also to the funding which is available -- which is made available to the participants through this SBIR program.

And the contest itself is formalized as the setting where, you know, the R&D is going to eventually produce an invention associated with some surplus, and the surplus is separated into the value and the cost of delivery of this invention, of this product. And, so, these two things are unknown ex ante and then they are sequentially revealed through the R&D process. So they are sequentially uncovered during the process.

So the model is going to assume that the success of the invention and the cost of the delivery, they are stochastical and monotone in the investment. And that the kind of the contest will result in winning if the invention -- the -- you know, the eventual invention is associated with positive surplus, meaning that the value is greater than the cost.

Methodologically, to actually link the data to the model and to uncover components of the model, the paper is going to assume -- the author is going to assume that the investment, which, of course, is not -- is not given in the data explicitly, so he will have to assume the investment is given to SBIR payment. And, also, he assumes that investment is monotone in value, and for some components of the model, he will have to assume that investment is actually -- investment equal to payment is actually optimal for the -- you know, given the surplus and given the environment.

So there are many good things that I can say about this paper. So, first of all, of course, it's a very timely effort thinking about this -- the optimal structure, the optimal features of R&D contests. You know, of course, these contests have been around for a long time. We know that [indiscernible] construction always involves a contest -- always involves a contest stage where people compete, where their, you know, multiple designs compete and then whoever proposes the best design gets to supervise the construction.

So these contests is seen then before, but we see more and more of them recently where the Government or private firms run contests to choose the best design or to kind of to generate more innovation in a particular area so one kind of very prominent example that I'm sure a lot of people heard about is this hyperloop pod competition which was run by Musk and Tesla company.

So, yeah, so this -- there seem to be a lot of interest in these contests recently, especially in the private market. So, again, timely effort.

So what else? So first of all, I have to commend the author for making, like, I am sure a pretty substantial effort of collecting the data that would be informative about this environment. So he

clearly had to go to multiple sources, put it together, to make it, you know, informative and coherent.

Second, you know, despite his best efforts, the data was limited. You know, there are multiple shortcomings in this data. And, so, he made a pretty substantial effort to design a model that is going to take a maximum advantage of the data that are available to him. So it's also, you know, a nontrivial -- nontrivial contribution here.

Also, you know, given the model, he proposes a new identification strategy, which very nicely takes advantage, leverages the features of the bargaining features and also the selection into the -- into the third stage that he has in his model. It's a very nice identification strategy. People probably will want to use it in the future.

Again, the paper provides a number of insights into how these contests should be optimally designed. I perhaps should not spend too much time going into it because I do want to mention a few -- a few kind of concerns that I had when reading the paper. So my main -- you know, a number of concerns that I have are related to the measurements of things in the paper. So first of all, this whole concept of

value surplus/profitability of the invention.

So to the best of my ability, the way he measures -- you know, to the best of my ability to understand what was written in the paper, the way he measures it, he basically links a particular R&D contest to the subsequent acquisitions by the Department of Defense. And, so, basically, the surplus is measured by the observed purchases by DOD of the invention which came out of this contest.

Right, so, first just a purely technical comment. It wasn't immediately clear to me whether the way he thinks about profitability was a per-unit profitability or kind of sort of lifelong, overall profitability. On one hand, the model seemed to be talking about per-unit profitability because we talk about the model looks at this cost of delivery, a unit of the product, right? So it's kind of a per-unit profitability.

On the other hand, what we measure in the data seemed to be more for multiple unit, right, in these profitability, and you would think that this lifelong profitability would be the right thing to take into account when thinking about investment, right, because that's what they anticipate to be the return to the -- to the R&D process.

33

And here I anticipate at least one concern. So even restricting kind of return to the investment to purely DOD acquisitions, you know, the firm has to forecast the future demand for the invention. So it's one thing to be able to assess how much they can extract -- you know, what would be the value of a given unit of the product. And it's an entirely different forecast and the procedure thinking about how many of those units they will be able to sell to DOD. So that's one thing.

The second thing is you may not be able to see the full realization of the demand in your data because, you know, so maybe they bought right now a few units, but maybe more purchases are coming in the future. You don't know, so not the full demand is perhaps realized in the data.

Second, when I was looking through the documents on the SBIR website, they keep emphasizing this, that all these kind of inventions, they have their profitability but then -- but then show it is not necessarily restricted to military uses. They keep encouraging the participants to seek kind of private sector sort of applicability of their inventions. And they keep emphasizing that the invention may be useful also as a stepping stone for

34

the future products that will be developed, right? So they're all like a broader sort of surplus coming out of this invention.

And, so, perhaps some of it you cannot -- you know, cannot measure just like obviously there are limitations to what is feasible, but I would be, first of all, a little bit concerned about the private sector potential, right? And one way to deal with it, which is perhaps not ideal but maybe better than what is done right now, is that maybe limit your attention to purely military so that pure -- to the project that are aimed at very clearly military uses.

So for example, these game-related projects, right, the virtual-reality-related project, for sure they will be placed on the private market as well, and so this is something that you cannot see in the data, and it perhaps is a shortcoming of the -- you know, like it's a pretty serious distortion in your measurement of the value. Okay?

So why should we worry about this? So, one -- one thing is VM is measured in surplus. VM is measured in the social surplus, and that is what you aim to maximize by your design of the contest. And, so, obviously this is -- this is already not ideal, but all -- you know, if you mismeasure social surplus,

35

if you define it incorrectly, then you probably are not going to be able to correctly predict the optimal investment that the firm would want to make, you know, when doing this R&D process, right? So that also is going to induce the distortion in your analysis of the investment.

So in general, you know, I understand, the data are limited. You do the best you can, but I would think a little bit more if there is anything you can do, like, additional about the investment because, you know, even -- even on this SBIR website, they do say that they provide seed money, right, so which already kind of says that it's probably not equal to the investment, or at least maybe again selection of projects which it's more likely to be exactly equal to investment. It's a little bit -- you know, at least acknowledge it in the paper so that people are aware that the results are subject to this, you know, possible shortcoming.

Okay, so, another concern that sort of a little bit nagged at me when I was reading the paper is whether you measure competitors correctly, right? So it seems that the way you think about competitors, you always think about people who are participating in the same SBIR contest, right? But the SBIR only

36

finances R&D by small businesses, right? And, so, potentially, there are other businesses out there that are doing similar research and, you know, the SBIR companies may not be aware of those competitors and are taking them into account when making their investment decisions.

So, again, why does it matter? Well, it matters, first of all, for your bargaining stage because that is going to influence the Government's threat point, and as I said, it may matter for the optimal investment.

So, again, what can be done? I understand. So one way to do -- to deal with this, again, if you reduce -- reduce the set of projects to those that look specifically at the military uses, perhaps you can go to this -- go back to this DOD database that you used and look at the SBIR topics, related acquisitions, which involve non-SBIR firms, right? So that could help you to define the set of other potential competitors, so other firms that worked on similar topics and kind of eventually got a scoop, like kind of beat the SBIR companies. So at least -- at least one way to address it.

So another concern which perhaps is not -- is of a smaller magnitude but nevertheless may be worth acknowledging in the paper, so this concern is about whether we are able to recover the correct distribution of values in this analysis. So one concern that I had is I already know that we only see the small companies in the data, but another thing is that the SBIR participation does involve some restrictions -- or does impose some restrictions.

For example, the products that come out of the research, which is based on SBIR financing, they cannot be exported. They cannot be sold abroad. Also, any patent that comes out of SB-funding research has -- so the Government has the right of free licensing of this -- of this patent for any future production.

So clearly firms are going to take this into account when deciding whether to apply -- to even apply for SB funding, right? So you would anticipate that there will be some selection and therefore we are not going to see the full distribution of values perhaps, you know, on the basis of the data that we -- that we use -- that we see in the -- you know, coming out of SBIR program.

So, again, even less a concern that something nevertheless that is worth acknowledging is that the social surplus generated by the participants

in this SBIR program, it is larger than just the explicit profitability, the explicit profit that they collect by selling, you know, the product that was eventually developed in this contest. So specifically, some of the losing ideas, they still result in the published knowledge. They result in patents. And they serve as a kind of fodder for the future research, and then in this way they contribute to social surplus.

But, also, the guys who lost, you know, the SB contest, they nevertheless may sell their invention in the private market, yeah, okay, so they will sell the invention in the private market so that, again, there is some surplus generated here which is perhaps not taken into account.

Just one thing to say, again, the data, very limited, we are doing the best we can given the data. But if in the future we have more data, like one concern that I had is that perhaps we do not get to learn the value before we learn the cost. Perhaps this process happens kind of concurrently, and in the first stage you only get a signal that is then [indiscernible] clarified in the subsequent stages, and so then identification will have to be adjusted accordingly to have such a rich environment.

And then the final comment, even of smaller sort of importance is that the heterogeneity of the projects clearly is important here. And right now, what is done in the paper is that the author allows for the distribution of value, distribution of cost, and, you know, the payments by SBIR to be sort of scaled in the same way, right, by exactly the same sort of factor. You may think it's too strong, right, it's a strong assumption.

Perhaps it's the best we can do right now, like, given the data, but, again, perhaps it's worth thinking that maybe the scale for values, the scale for costs could be different, in which case we may need another variable, we need another measurement. One thing that I thought, you do get in these proposals there is an estimate of the cost that they provide in the second stage. So perhaps that's what you can use again as a variable to help you to kind of to better capture the scale of these -- of these inventions -- of the distribution of cost for these inventions.

So this is all I have. Thank you very much for your patience. And great paper. I hope to see more research in this area.

MR. RAVAL: All right, we have time for one

question.

MS. JIN: A very interesting paper. I'm wondering to what extent you observed repeated interactions between the small firms and DOD. If a firm got rewarded in phase three, DOD potentially would have much more information about a firm when they deliver phase three, and would that information sort of help them in the future project selection?

MR. BHATTACHARYA: Yeah, so these firms, the conditional winning of phase two contract, I think there the data said, like, four or five times on average. So there are repeat purchases. There's some DOD history. I haven't looked at winning a phase three contract and whether you win more phase two contracts, but there's probably an effect.

Winning one phase two contract in the past tends to increase your probability of winning phase two contracts, but after that, it's pretty flat. So I think there's some learning between the DOD and the firm, but -- but maybe not for an excessively long period of time.

Okay, is that it?

(Applause.)

MR. RAVAL: All right, next we have Allan Collard-Wexler from Duke, and he's going to present

Market Power and Product (Mis)Allocation in OPEC: A Study of the World Oil Market.

MR. COLLARD-WEXLER: Okay, so we came at this project by thinking about misallocation of production, so inputs going to the wrong firms. And we wanted to understand what was the effect of market power in generating some of this misallocation of production. And the setting that we've decided to look at is the world oil market. And this is a market, I think, that's very interesting to study for productive misallocation.

First of all, there's a large cartel that's been active for a very long time, OPEC. It's a homogenous product market where we're going to be able to kind of understand production costs at different parts of the world in a kind of very comparable way. And I think we're interested kind of in the effects of market power, but we haven't spent a lot of time on kind of cost effects of market power and I think this is where we're getting at.

And, finally, I think it's also to bring questions of market power to the kind of misallocation literature. And that's why we started this.

So this is going to be the main graph to explain to you the distortion and what we're trying to

measure. So imagine that -- I don't know if there's a laser pointer? No. So imagine that you have a cartel with marginal cost 1, so they're the low-cost guys, and the socially efficient thing would be for marginal cost 1 to produce everything. But this low-cost producer happens to be a cartel, and they're restricting production to q1. And because they restrict production to q1, you've got other producers in the world, and those are represented by this marginal cost, f, that's increasing the jump-in. And there are going to be some competitive fringe. They produce all the way until marginal cost equals to price.

Now, the typical thing that we do when we looking at this is look at the quantity distortions, look at the Harberger triangle that we're -- we're not producing the socially efficient amount; we're producing less; and that's causing a welfare loss.

And what I want to draw your attention to is there's also another loss in this setting, and that's that production's being allocated to the wrong people. So even if you wanted to produce q rather than the social quantity, you wouldn't produce q that way efficiently. And, so, this trapezoid that we shaded in is just representing the increase in total cost of

production, which is a welfare loss, because the cartel's leading to inefficient allocation or production between inside the cartel and outside the cartel.

We call that -- that shaded trapezoid productive distortion, so it's a distortion that affects the -- that affects the cost of production. And the goal of this paper is to try to measure how big that distortion could be in the context of OPEC.

Now, as soon as you start this, there's a problem with oil which is it's a renewable -- it's a non-renewable resource. So, you know, there's this question of, well, if I don't use this field today, I can -- I can just use it tomorrow. And, so, what we're going to do is to take this kind of depletable resource setting seriously and that welfare gains are going to come from we should be producing at low cost -- we should be kind of moving low-cost fields kind of early in the production order rather than later because -- just because of discounting, it's going to be more efficient to use cheap resources before you use expensive resources.

So it will just be a dynamic version of that productive misallocation graph that I just showed you. So really it will be all about the timing of

extraction to take that depletable resource context seriously.

There's been some literature that's tried to get this productive distortion measure, and the one I want to point out is in the electricity market, Borenstein, Bushnell, and Wolak have something similar because electricity, there's inelastic demand curve, so there's no quantity distortion, so all you're left with is productive distortions. And there's a large literature on misallocation and on cartels and we're really trying to join the two together. And there's less literature on OPEC than you think, so we're also trying to add in on that.

So what we'll find is that over the period 1970 to 2015, cost of world oil production are 10 percent higher than they ought to be because of the OPEC cartel. And this productive distortion has -- over this time period has a welfare of $163 billion. So it's saying that these productive distortions could lead to welfare losses due to market power that are as large as anything that's been documented. And that's why we should think about them when thinking about the welfare impacts of market power.

Okay, so some background on oil is there's large cost differences between oil producers. I think the 90/10 -- the 90th versus the 10th percentile have like a nine-to-one difference in cost. For manufacturing, that's like three to one. And it's pretty easy to understand why there's such dispersion of costs of oil extraction. You know, this is in West Texas. These are just stripper wells, so this technology is -- you know, you could have done this 70 years ago. It's reasonably easy to do.

And then here's another oilfield. This is off the North Sea off Norway, and this is like building a skyscraper in the middle of the ocean. So there's kind of natural reasons why different oil deposits are going to have very different costs. And, you know, and that's why which oilfield gets extracted when is going to have kind of meaningful effects on total costs of oil production.

OPEC is these countries. I would say when you read about OPEC, it's an imperfect cartel. So they use quota arrangements rather than, say, telling Saudi Arabia you produce everything and send a check back to Gabon. So there's no transfers in this cartel. There's instances of cheating on quotas. A lot of people would think that a -- the market power of OPEC is just unilateral market power by, say, Saudi Arabia and Kuwait, so it might not even be a cartel

the way we model it but just a set of leading countries that -- that exercise unilateral market power.

Why is this important? It's -- you shouldn't even expect OPEC to basically minimize costs within the OPEC cartel given their -- they're an imperfect cartel mechanism.

And, you know, OPEC is about 40 percent of world oil production, and outside of the OPEC is the rest, and some countries like Saudi Arabia and the United States are 10 to 15 percent of global oil production.

And this is just another way to say it, the dotted line here is the OPEC market share, and I've overlaid the price of oil over time, and there's these big instances in '73, '81, and then recently where the price of oil has these large spikes. So there's a lot of movement in the price of oil generated according to observers by OPEC's decisions to cut production. And that's why production starts to spike -- costs -- prices start to spike.

So to understand misallocation of production, we needed data on many oilfields and their costs and their production. And we got this from a Norwegian energy firm called Rystad Energy. There's

about three different firms that do micro models of the world oil market, and they assign, you know, costs and reserves and production to basically all the -- all the oilfields on the planet. And we're leveraging this data on 13,000 fields.

The data is going to be at the field level. So, like, Ghawar Uthmaniyah is one of the world's largest oilfields in Saudi Arabia, and that has, say, like 800 rigs on it or more. So some of these fields are -- it's not a single well. It's a field with, say, up to thousands of wells on them. So that's the level of the data. And this is just to say that we have detailed rich data on these individual fields, such as reserves or when they were discovered or how much they produced from '70 onwards.

The first thing you might think is, well, maybe -- maybe the reason that OPEC produces, say, 40 percent -- 30 or 40 percent of the world's oil is because it's limited on reserves, so there's only so much oil in the ground that it has and that's what's constraining it. And just as a first pass, you know, OPEC might be 40 percent of production, but it's about 50 percent of reserves in the world. And if you do something simple, which is to say, like, what's the ratio of reserves to annual production, so non-OPEC

with current reserves, they can produce for the next ten years. In OPEC, the same answer is 19 years at current production.

So it's not just -- when I say that OPEC is producing too little, you can see that because they're exploiting the reserves less intensely than non-OPEC members. And that's something we noticed.

So I want to give you an idea of what the variants in costs looks like across the world, both OPEC and non-OPEC members. And what we've done here is we've plotted what I call, like, annual costs, like costs in a year divided by production in a year, and the black bars are the 5th to the 95th percentile across all the oilfields in that particular country.

And just for -- just to benchmark things, I've also put what the price of oil is on top of that, so that gives you an idea of, you know, how these costs compare to prices. So this is Saudi Arabia. You know, we estimate today they have a cost of, say, $10 a barrel. So they have the cheapest oil reserves in the world. If you look at countries like Nigeria, which is an OPEC but isn't one of the countries that exerts unilateral market power or punishes cheating, they have costs that are quite a bit higher, say a mediant of on the order of $30 a barrel recently.

A country like Russia has costs of about $20 a barrel. They're not in OPEC, and you might wonder, you know, why aren't they producing more, given they have low costs. And here I'll just note, and this is outside of the discussion for now, they have a 50 percent pipeline tax. So basically half their revenue is a direct royalty to the government. So you might expect there's other reasons why marginal cost and prices might diverge in terms of production choices.

And then if you go to the United States, the 90th percentile well in 2014 had costs of well over $90 a barrel. These are mostly fracking, by the way. And this is exactly the kind of productive distortion that I want to get at, which is there's tons of oil in Saudi Arabia in these $10-a-barrel fields. And the price kind of skyrocketed to $90 a barrel in 2014. And then there's a question of why didn't Saudi Arabia produce more, and, well, that's because, you know, that's how it exerts market power by holding down production. But then when the price is $90 a barrel, things like people fracking in North Dakota at $90 or $80 a barrel, like really expensive oil production starts to enter the market. And that's the kind of productive misallocation I'm talking about.

And you see the same pattern in Canada

where, you know, some of the most expensive oil production that you can see in the world, which is exploiting tar sand, starts turning on in the 2000s. And, again, given that there is available oil that's at $10 a barrel, squeezing oil out of -- out of the sand at $100 a barrel seems like a very inefficient thing to do.

Okay, so, this is kind of the evidence for just how big the cost differences are across countries. But I wanted to get at this measuring the productive distortion, so I wanted to estimate that shaded trapezoid. And, so, this is what we're going to do. We're going to propose a measure -- a definition of productive distortion, and that's the difference between -- remember, this is dynamic, so it's going to be the net present value of the realized cost of production, given what we see, which we'll assume is due to the activities of the OPEC cartel, versus the net present value if firms took prices as exogenous so that -- that means they're acting as if they were in a competitive world, but there's an important caveat, which is they took prices that are exogenous, but the total production of oil in the world will be the same as what was realized in the data.

Why am I putting that caveat that total production is the same as what was realized in the data? It's just that I want to keep total production in this counterfactual at q, what actually happened as a net production distortion figure. So we're not going to be kind of playing around with increases in total aggregate quantity, just holding quantities fixed in the competitive counterfactual versus the data and just looking at how the allocation of production differs in those two worlds.

Okay, we need to put a few assumptions. We're going to have a very long run view on costs, so like cost of developing an oilfield from nothing until depleting the field, and that's going to mix startup costs, fixed costs, marginal costs and so on, and I just want you to think that over a long time period you can kind of combine these together into a single kind of unit cost.

In the paper, we do some derivations with production functions to get something that looks like CFT, which is just a constant marginal cost. And there's going to be some work here. This Mu-st factor is just going to try to pick up that. It turns out, like, the costs of renting a rig move a lot from year to year, like when the price of oil is $80, it costs

three times more to rent a rig than when it's $30 a barrel. And so this Mu-st thing is just trying to capture a variance in input prices of drilling, and that's why we need to have it in there.

And we're going to assume that this -- you can just think -- for the purposes of this talk, you can just think of those costs as just being constant over the whole time period. There's just a CF cost of a particular field.

Now, with this kind of linear marginal cost structure that's constant, you get a very nice characterization of the competitive equilibrium. So just as the competitive equilibrium firms are maximizing the NPV of profits, subject to a reserve constraint, in the paper, we have a way of characterizing this equilibrium, which is through what we call a sorting theorem, which is just the lowest cost guy starts producing all the way up until you've satisfied the total quantity, that restriction for that year, and then you move on to the next year, and then -- so you just keep depleting the cheap fields up until the quantity constraint and then you move on.

So really it's just saying the cheap guys go first. That's what the equilibrium will look like. And so this allows us to kind of very simply solve for

53

a competitive equilibrium or 35 years with 13,000 wells without having to do too much to it.

So I want to show you results from comparing what happened to this competitive counterfactual. And I'll first start by telling you what happens if you just look at this competitive counterfactual applied to a single year. And then I'll do what I call a dynamic version, which is I'll look at the competitive counterfactual but from -- starting in 1970 to 2014. And those differ because oil is depletable, right? So if you extract something in 1970, you can't use it in '75. So that's why the dynamics are different. The reserves are kind of the state variable here.

Okay. And there's a number of modeling assumptions like discount rates, how quickly you can extract oil from the ground, given the size of reserve. What do you assume about the discovery process of new fields. So we have to make assumptions on that. It makes sense to kind of run this competitive simulation all the way until every single drop of oil in the world has been depleted. So really the only difference between competition and market power is going to be just the timing of oil extraction, not is every reserve going to be extracted. That will happen, so we're just going to

54

have to assume things like after 2015 the entire world reverts to competition, just to fill out kind of the later years of the model.

Okay, and then there's some work on estimating costs, which I won't get into, but is done in the paper. So I just want to show you in 2014 what happens to output shares. In actuality, that's the left side. So like the Persian Gulf OPEC members produced 26 percent of the world oil in 2014. In the competitive counterfactual, they would have produced 75 percent. That's not surprising. Most of the world's cheap oil is in the Persian Gulf, and so in a competitive world, you just see production kind of ramp out from there.

Interestingly enough, the members of OPEC that are not in the Persian Gulf would actually see their shares drop a lot. And if you wondered, well, you know, is Venezuela really doing anything for OPEC? You know, our answer would be it doesn't look like it. It looks like it's producing more in OPEC than it would in a competitive world.

And then, of course, if it's going to the Persian Gulf, it's coming out from somewhere, and it's coming out mainly from non-OPEC members, so the U.S. produces 13 percent of the world's oil and it produced

55

1 percent in the competitive counterfactual. So, again, this is just showing you how market shares would change under our kind of competitive system versus what we have.

We then wanted to say, well, what is the actual welfare cost of this different allocation of production. And, again, this is for just one single year, and, you know, the competitive cost of production, which I'm going to call optimal for second welfare theorem reasons, is $121 billion, whereas what we see given the actual allocation of production in 2014 is a cost of $240 billion. So there's basically the cost of oil is about twice as high as it would be in our competitive counterfactual.

Now, in subsets, that doubling of costs is really strong, and I'll explain to you why. If you look at, say, if I just fixed production levels within each country, but then I made things competitive in each country, so I allowed all the fields in the country to produce in a competitive way, then costs would be 203 billion. So it would be a $40 billion savings. So that's like saying there's inefficiency within the country that we're picking up.

And, now, attributing that kind of inefficiency to market power seems like it's a very

56

odd thing to do. It could be measurement problems. It could be errors by the producers. It could be expectations that weren't realized. So it could be a lot of other stuff.

So we're going to try to do something conservative and say -- and this is in the optimal OPEC quantity -- which is if you just fixed in this competitive counterfactual, the total -- that 40 percent of the world's oil comes from OPEC and 60 percent of the world's oil comes from outside of OPEC, how much more expensive is the total cost of production than the fully competitive counterfactual. And that's what this $154 billion number is. It's saying just the allocation between OPEC and non-OPEC countries is causing the cost of oil to be $33 billion higher than it would be without that, that restriction on where all oil comes from.

And, so, we're going to use that kind of number to kind of -- we're going to call that the effective market power here. And, you know, we have these distortions plotted over time. They get bigger as you get -- in the periods when there is spiking price of oil, which shouldn't be surprising. You know, it's not when the price of oil is $30 a barrel that you expect big misallocation. It's when the

57

price of oil is $90 a barrel that you expect very expensive oil to kind of hit the market. And that's a productive inefficiency.

Same exercise now. Instead of just doing it for one year, we're basically simulating out kind of what would happen over time. That's our dynamic counterfactual. And you get these results that, like, in the 1970s, OPEC should have produced in a competitive world 90 percent of the world's oil. It's even stronger than that.

There's, like, three fields in the world -- Ghawar, Greater Burgan, and -- sorry, there's two subfields of Ghawar in there that would have basically produced everything. So they're the cheapest oilfields in the world. They're like at $5 a barrel. In the competitive counterfactual, you should just deplete them immediately. And then once you've depleted them by, you know, 1990 or so, you let other producers kind of kick in. But really it's just saying the ordering of those fields is very strange. They should have been depleted immediately.

And then we do kind of the same kind of costs but for this entire path from '70 all the way on to 2014 or all the way until 2100, which just represents until all the world's oil gets depleted.

58

And you get things like the actual cost of oil was 2.1 trillion. The competitive counterfactual would have been $1.2 trillion. So, again, the same order of magnitudes.

And then if you look at, you know, how much of that is because of -- of that increase is because of that $900 billion number is because OPEC and non-OPEC market shares are -- the market share of OPEC's being fixed at what it actually was, the answer is 148 billion. If you look at what's coming from across OPEC member distortions, that's 85 billion. So this is just accounting for where these -- where this misallocation is coming from.

So the headline numbers we're going to bring up here are if you just count the fact that just constraining OPEC's market share to be what it was in the data, you get a number like 148 billion. If you count not only constraining OPEC's market share but also that within OPEC members production is being misallocated and there's good reasons to believe that, like, Venezuela is producing too much and Saudi Arabia too little within OPEC because of how the cartel is organized. They don't have transfers, so they use kind of market share to move things around. Then you get a number of 233 billion, you know, incorporating

59

those within OPEC distortions.

Let me add one more twist to all of this, which is, you know, we're kind of worried that there's -- we're looking at actuality versus a competitive model where there aren't any other distortions. And you might think, well, what we should really be comparing this is like to a second-best theorem where, you know, it's competitive but there's also other distortions like distortionary taxes. So we know that there's a lot of distortionary royalties here.

So even if you move competition but you keep those distortionary royalties, you know, you're not -- you're not going to get cost-minimizing production. Or maybe there's all sorts of other wedges that might be distorting production that don't have to do with market power but would still affect the competitive counterfactual.

And in the paper, what we show is that even if you kind of condition on what's the effect of market power with those distortionary taxes or with any other distortionary wedge that causes, you know, the low-cost fields, say within a country, not to produce first, we get -- the OPEC numbers that we've been measuring turn out to be very stable. So it's adding these other distortions. So we feel reasonably

60

confident that even in the presence of some other distortions these effects of OPEC seem to persist.

So I'll just conclude. There's countries with clear market power. They're in the Gulf. They have very low cost of oil production. If you push production towards those countries, like you would in a competitive world, cost of oil production would drop substantially, and this leads to enormous welfare effects due to market power, and again, not the traditional channel, but the quantities are too low, but instead that the allocation of production seems to be distorted by the cartel.

Okay, that's it.

(Applause.)

MR. RAVAL: So we have Hugo Hopenhayn from UCLA to discuss the paper.

MR. HOPENHAYN: My discussion will not take all the time. This is a great paper. I mean, there's a large growing literature on misallocation, sort of more in the macrodevelopment side. Allan and the coathors here have identified perhaps the -- if you look at that literature, one of the big failures of the literature is that while measured misallocation tends to be very large, identifying the causes of that misallocation has been, you know, really poor.

61

And policies that we usually think about, you know, generating distortions in allocation, distorted taxes, subsidies to firms of one type or the other, explain, you know, a very small fraction of misallocation. For example, in Hsieh and Klenow's study of misallocation in China, all the observables that you have at first had explained about 10 percent of the misallocation. So finding from that perspective sort of there's a holy grail of finding, okay, well, what is behind this misallocation. And as, you know, previous paper by Allan and coauthor -- the same coathors have argued, I mean, a lot of this misallocation could be -- basically be it's backed out, you know, from structural specifications, misspecification of production functions, for example.

So this one is, I think, very valuable in that context because it brings in real data and a very clear reason for having misallocation. I'm not going to comment too much on the results themselves, I mean, in terms of data. It's not my strong point, those of you that know me. And the other thing is that I think, you know, the paper is very carefully done. There's a lot of issues that, you know, they had to make assumptions about how much you can extract, at what rate, you know, it's a maximum, this 10 percent

62

rate that you can extract in establishing their counterfactuals.

One thing that I would note only in terms of quantitative analysis is that when they look at misallocation within countries, in particular within countries outside of OPEC, there's very large misallocation. And in some ways, it's hard to think, I mean, that it's imperfect competition that is generating that in that, you know, it's -- the rest of the world is a fairly competitive market. It's not very concentrated production, and so it's very dispersed. And if you take price as given, you know, if you have a bunch of competitive firms, they would minimize cost subject to that price.

And, so, what is creating that, you know, big distortion that they find there, and then whether that could be used in some ways to get a sense of the extent of measurement error that there might be, and, you know, that, you know, taken to the other calculations, you know, sort of to as in Hsieh and Klenow they do and sort of to -- you know, make the values more, you know, relative to that normal measurement error, let's say, so I think that's, you know, the only caveat that I would want to point out. I mean, I still think, you know, the paper has a big

63

contribution, and these numbers come out very strongly and high.

So my question is more as to, okay, so we're at the FTC. And, so, you're concerned maybe about collusion and the cost of collusion. And, so, what is it -- from that perspective, what is the right benchmark? So -- and Allan very clearly pointed, the paper is about misallocation. And, so, the reference point in a paper of misallocation would be the optimal allocation, or what he called the competitive, which would be the optimal allocation a social planner would choose, which is the cost-minimizing or sort of present value cost-minimizing allocation.

But if we think about collusion and we're thinking about damages, and by the way, I think the paper is pointing to something that I don't know how aware people are, you know, how important it is in measurements, which is that we're used to these Harberger triangles, which are about, you know, a welfare loss is from cutting output. That's the triangle, but we're not used to this calculation that, you know, when the rate is -- potentially when there is, you know, this missed -- imperfect competition or collusion in this case, that generates another triangle, rectangle, or I guess the average would be a

64

trapezoid, no, of the two?

So I think their paper points out to something that is very important and that we should be more aware of. And I even thought about, you know, having more a macro perspective that what not to recalculate, you know, the old, you know, Harberger welfare triangles and add to those, you know, the rectangles that come from, let's say, a imperfect competition that as we know, those of you that, you know, have worked with Cournot models, that there is, you know, inefficiency in allocation as firms with different marginal costs produce.

Okay, so that -- this is what I'm going to say is that I want to put a little bit of this in perspective and, you know, ask, you know, what -- and sort of what is the right benchmark. And if we're thinking about the FTC's thinking about collusion, then banning collusion or eliminating collusion is not going to eliminate misallocation. It's going to give us the misallocation that is generated by imperfect competition.

And as Allan pointed out, I mean, the -- it's important that collusion here in the cartel is imperfect because we know that perfect collusion, I mean, with transfers actually could improve misallocation by actually having the output assigned efficiently and then doing the appropriate transfers. So it's not obvious that collusion, per se, hurts misallocation. Here, the source is going to be that collusion by raising the price will allow, you know, certain producers to come into the market that are inefficient and that would be not even producing in the absence of collusion.

So that's a question, what's the benchmark, and so I think that's kind of an important question that we should ask, and the question is -- I mean depends on what we're after. If we're after, you know, damages of collusion, there's one thing. If we're after misallocation and understanding differences in TFP, I mean, this is the competitive or optimal one is the natural one.

The other question that I'm going to ask is -- and I guess I played with this a little bit, I mean, there's really not a lot of -- I mean, there's really, except for that graphic you saw, there's no more theory in the present value allocation, there's no more theory in the paper. So I started playing a little bit and saying, okay, maybe theory will take me somewhere, and I'm just going to tell you where I got. I mean, that's -- and you'll see the theory is, you

know, first year undergraduate micro theory. So, I mean, it's not very fancy.

And, so, I'm going to address the suppressance of this high-cost fringe, make it worse or -- and we saw an expansion throughout time -- this high-cost fringe. Did that make things worse or better? I mean, in some ways, they -- you'd say that's good, I mean, they created output that would not produce instead, but, you know, they also contributed to this misallocation. Is that good or bad? That's what I'm...

So this is the -- I thought I was going to get some -- oh, how do I go back?

You've seen already these numbers. I mean, he didn't have them -- oh, yeah, he had the same table. I'm not going to comment any more. The lower numbers are the ones that correspond more to the exercise when the benchmark is -- or partly the exercise when the benchmark is, you know, eliminating the misallocation that -- the lower one between OPEC and not-OPEC countries. Of course, this is much smaller than, you know, the upper numbers, but this is the number that I think you know realistically you want to point out when the bench -- I mean, if you're thinking of this benchmark.

The other thing is that you are not going to be eliminating the whole misallocation within OPEC, and not even across OPEC because there is going to be in imperfect competition some misallocation -- residual misallocation. So these numbers would be possibly considerably smaller when you consider as a benchmark in let's say a Cournot equilibrium as opposed to considering as a benchmark perfect competition.

I don't know how easy it is to think about even doing that kind of exercise but, you know, I think it would be nice to have some idea of orders of magnitude, perhaps comparing for episodes where there was break in the collusion and sort of thinking of that perhaps as the allocations that you would see in the absence of collusion. I really don't know, I mean, what would be a good...

So this is what we know in terms of -- I mean, Cournot model, the markup rule, the markup of firms are proportionate to the market shares. From this, you can back out that there is residual -- I mean, that there will be a coexistence of firms with different marginal costs within some range.

I did some -- just to give you a benchmark for this, I played around with the linear demand model

and constant marginal cost. It turns out that the maximum misallocation that you can get there is if there is a single other firm with high -- you know, higher marginal cost producing. And the max is about the size of the trapezoid compared to, let's say, the welfare triangle is half that size. So the trapezoid is half the size of the welfare triangle, but it's already an important -- it already says that, you know, we are in that -- in that model. I mean, it's a bound. We would be missing 50 percent of misallocation just by looking at welfare triangles and not looking at misallocation.

Yeah, so, to talk about the counterfactuals, so here is a picture -- I mean a more stylized picture, a triangle instead of a trapezoid, but the same thing as what Allan presented. You know, we started the marginal cost, C. That's a quantity under collusion. The q, the small q, corresponds to the cartel's quantity, here assuming that the cartel has the same marginal costs as in his picture. This is just for expositional purposes. And the total quantity with collusion is adding the supply function that is depicted here, which would be the fringe firms, all of which have a marginal cost above the collusive -- or the marginal costs of the collusing firm.

And that would be, you know, this misallocation if you want a productive inefficiency, the CL, the -- compared to a deadweight loss. Now, if we, let's say, eliminated the cartel and as a consequence now we produce let's say a Cournot quantity, I mean, I don't know where the Cournot quantity is, you know, but if it proves a Cournot quantity, now you can see, I mean, that still there's going to be room for some fringe.

And, so, yes, there is a reduction in this triangle, but it doesn't disappear. And then, I mean, my picture, I mean, maybe suggests that that reduction might not be so large, unless it were a really large increase in the output, and the marginal costs of these firms -- like the supply function were concentrated in the upper levels, closer to price.

Second question here, so this says two things at the same time. I mean, obviously we're better off that there was this fringe, expensive fringe, because even though they introduced a misallocation cost, I mean, they're producing at a cost that is less than marginal cost -- sorry, marginal cost that is less than price. So they are contributing to welfare. And, so, yes, I mean, it's

good that we have these firms; however, you know, it's not -- you know, we need to take into account that triangle, so -- in terms of the implications of cutting output by OPEC.

So the fringe expanded considerably during this period. And, in fact, you can ask the question, did that fringe expansion hurt welfare. And, actually, it can, and this is sort of -- maybe it's a little bit nerdy, but, you know, I just want to show you because I think it's a very nice calculation that you can make in this respect. And this is sort of a very -- again, going -- you know, a micro-level one.

So think about just to explain this picture, c here is the marginal cost of the cartel. Let's say here I'm taking C-zero to be the marginal cost. I'm thinking in cost to make it simple of a fringe. Initially, the capacity of the fringe is Q0. And I'm going to consider an expansion of the capacity of the fringe. And let's say that given that that's a capacity of Q0 of the fringe, now we think of the cartel, it's going to be best responding as a cartel to this capacity of the fringe, and its best response that say that total output would be Output-Q. Okay, so this is sort of the initial equilibrium with a given capacity at Q0 of the fringe.

Now suppose that, you know, you start doing fracking in North Dakota and all these things that expands the capacity of this fringe, okay? So -- and by the way, obviously if we go back, we have the two sources -- the two -- here's the rectangle -- yeah, sorry. Here's a rectangle. The size of the rectangle, CL, versus the deadweight loss. I mean, that gives you sort of what are these two components of welfare losses.

And, so, now, let's say you got the expansion of the capacity of the fringe, okay, so we're going to get a little extra -- two little extra effects. One is positive, the right one, that is. We're going to decrease deadweight loss, and that's kind of that sort of trapezoid but let's say approximately a rectangle up there.

And we have this other rectangle, which is the increasing in this misallocation cost. Okay? So which is bigger? And this will tell us whether that expansion in the fringe is something that hurt or actually improved welfare, okay? And, so, it's easy to see here that the -- this isn't a Cournot or linear sort of -- I'm assuming linear demand. So what do we know about linear demand, that when one -- you know, when output is expanded, the response of the cartel

would be -- here's like two firms, is to cut by half of that expansion its own output, not by the full size, you know. So there's an accommodating part and so total output is really -- the total output expansion is only half of the expansion of the fringe.

And, so, these are essentially rectangles that differ in the base, okay? And potentially also differ in the height. One goes all the way to the demand function; the other one goes all the way to the cost, okay? So a calculation is quite simple. The change in CL is proportional to the difference between marginal cost of the fringe and marginal cost of the cartel. The change in Q is -- sorry, in deadweight loss is proportional to the difference between price and the marginal cost of the cartel, but divided by two because of the compensating effect of the response of the cartel.

So the total change is of the order of, you know, PC, divided by two, minus C -- 0 minus C. So here if the cost of the fringe is above half point between the marginal cost of the cartel in price, then this expansion of the fringe is actually bad for welfare. And I think this is kind of probably the case. You know, I mean, this fringe was -- you know, their costs were, you know, way above -- much closer to price than they were to marginal cost. So the answer there would be expansion of the fringe was costly, and it actually hurt the welfare of the world, even though the existence initially of the fringe is something that, you know, obviously possibly contributed to value.

So, I guess, you know, I don't have much more to say. I think it's a great paper. I think, you know, it's a very careful empirical analysis and dynamic modeling, in taking sort of the right way of approaching this problem. There are a lot of corners to cut, you know, that's inevitable, and I think they did a good job in that sense.

My main point is what is the correct benchmark if we're going to think about misallocation or versus we're going to think about collusion and trying to measure -- assess the damages of collusion. And, well, I think those are sort of the main points that I wanted to make.

MR. RAVAL: All right. Again, we have time for one question.

MR. RAMEZZANA: So I was wondering how general your optimal extraction path is, that is did you first start with the low cost field? Now, in a very stationary environment like that, I can see that,

but in a more complex environment, in which one could have, you know, big future shocks, unexpected shocks to the need for oil or the marginal utility of income on ability of a country to pay for oil, I can see why maybe you don't want to use all the cheap oil soon.

Now, if you can commit perfectly to an extraction path, yes, then you do that. We use the cheap ones soon and some better stuff happens, we'll take it. But there's no commitment in this world, so maybe like the United States or somebody else, you don't want to find yourself in 20 years really needing oil during a period of crisis and having high costs. So that was just my question, you know, relative to the welfare criteria.

MR. COLLARD-WEXLER: So the way I see this is like is Saudi Arabia not extracting everything now because it can't just put the money in a bank and then, you know, it's using the oil in the ground as some kind of commitment to savings? And, so, I'm sure that that kind of institutionally can these countries save that way, I'm sure that's an issue here. How big it is, I don't -- I don't know, but I think that's the -- but that's the gist of the question.

So thanks, Hugo, for the discussion.

(Applause.)

MR. RAVAL: All right, we have Jihye Jeon from Boston University that's going to talk about Learning and Investment under Demand Uncertainty.

MS. JEON: All right, so thank you so much for this opportunity to talk to you about my research. So I'm going to start off by saying that in many capital-intensive industries, firms experience large waves of investment. And firms in these industries also invest in long-lived capital while facing demand that's highly volatile. So their expectations about how demand will evolve in the future will likely play an important role.

So the container shipping industry provides an example of these boom and bust cycles of investment. So in the figure that you're looking at, the blue bars are quarterly investment in new ships, and the red line is the price of investment. And, so, you'll see that investment is highly volatile, first of all. Also, it is highly concentrated in times of high price of the investment.

So in this industry, firms are exposed to sharp swings in international trade demand, but at the same time, supply is hard to adjust in the short run because there's time to build and also because firms tend to stick to their preannounced schedules. What

happened recently is interesting. So there was a huge investment boom when trade demand was booming in the mid 2000s, and when demand collapsed after the financial crisis, this led to a huge amount of oversupply in the industry.

And, so, in this paper, I want to understand what drives these boom and bust cycles of investment and how firms invest under demand certainty, and I'm going to focus on the role of information. And I'm going to think about these things in a setting where there's market power and strategic considerations.

And, so, what I mean by focusing on the role of information is the following. So the standard way of thinking about agents' beliefs in a dynamic oligopoly model is to assume that firms know the true data-generating process. So in this environment, the only source of uncertainty would be about what exact amount of realization I'm going to receive today. Okay?

In addition to this type of uncertainty in this paper, I'm going to incorporate uncertainty about the demand process itself. So why do I think that this is an important factor? So a lot of industry experts were trying to understand what was going on and what drove this oversupply problem, and a lot of them attributed it to a firm's limited information. So this one particular quote says, "The industry extrapolated the good times and foresaw an unsustainable rise in demand."

There are also a growing body of studies that use learning models to describe agents' beliefs with respect to macroeconomic shocks and trade demand are highly correlated with these shocks. Lastly, of course, the benchmark rational expectations, assumptions are appropriate for many of the settings that we study; however, there are also many settings where this may not be the case. So firms may be new to the environment, for example, or the environment itself may be subject to some structural changes due to policy shocks or other exogenous shocks.

So in this paper, I'm going to try to address these questions. So first of all whether a model that incorporates learning about this aggregate demand process can help us understand the -- how firms are investing. And, also, how this learning in agents' beliefs interact with strategic incentives of the firms. Lastly, I'm going to think about whether the modeling choice of firms' expectations matter when we do policy evaluation or welfare analysis.

So here is the overview of my approach. I'm

going to first propose a dynamic oligopoly framework where agents are forming and revising expectations about the aggregate demand using information that's available to them at the moment they're making their decisions. Agents may believe that the process itself is changing over time, so the natural way to model agent beliefs in this case would be to allow them to put heavier weight on more recent observations. So I'm going to allow for this.

I'm also going to look at various other alternative models of firm beliefs. I'm going to compare predictions of my model to those of the other models. I'm going to estimate this model using firm-level data from the container shipping industry and then conduct counterfactuals with respect to combination, demand volatility, and scrapping subsidies.

So the first set of counterfactuals, which is with respect to competition and allowing coordination and investment, is going to highlight how strategic interaction plays a part in overcapacity as well. And I'm going to do this exercise under two different informational regimes, so under the learning model and the other one under full information, to look at this modeling choice would matter.

So, of course, one of the biggest challenges in thinking about selecting an appropriate information structure is that as researchers we do not observe agent beliefs directly. And as Manski points out, it's hard to identify information and model parameters simultaneously.

So the strategy I'm going to take in order to tackle this problem is the following. So, first of all, the standard approach in estimating this type of dynamic oligopoly model is to look for objects like investment costs, entry costs, or exit values that can rationalize firm behavior that we observe in the data while imposing the full information structure.

For my setting, I have data on shipbuilding prices, as well as prices on scrapping, so scrapping values basically. So I'm going to use this data directly and then instead I'll focus on identifying the model of firm beliefs. So taking the investment costs and the scrapping values in the data as given, patterns in the data such as investment volatility and the correlation of -- between investment and demand will tell us something about agent beliefs.

I'm also going to, as I said, consider various alternative models of firm beliefs. Of course, these two things are highly reliant on

structural assumptions that I make in various parts of the model. So as a more model-free way of thinking about this, I'm going to rely on GDP forecast data. And, so, the idea is the following: The GDP in the destination region is highly correlated with trade demand. And, so, the idea is that the correct model of firm belief should yield demand forecasts that are highly correlated with or consistent with the GDP forecast.

Okay. So I'm not going to be able to get into all the details of the model and the estimation, so let me just highlight some main findings that I have. So, first of all, I find that learning raises the volatility of investment and the correlation between demand and investment. And this is going to -- what's going to help me predict these boom and bust patterns that we see in the data.

In particular, I find that agents put heavier weights on more recent observations, such that the relative weight on an observation from ten years ago is around 45 percent compared to the weight on the most recent observation. And this is also confirmed with a validity test that I conduct using GDP forecast data.

I find that strategic incentives increase

81

both the level and the volatility of investment. And learning intensifies these forces. So, in summary, learning amplifies investment cycles, both through, first of all, leading agents to revise their beliefs as they face demand volatility, but also through intensifying the strategic incentives.

Lastly, I find that the modeling of firms' expectations has policy implications. So, in particular, the full information model underestimates welfare gains from a particular merger that I consider between the top two firms in this industry.

So the key ingredients in the data that I use is the following. I have route-level data on prices and quantity, and I have firm-level data on capital investment and deployment and the firm routes that they operate on, as well as some data on shipbuilding and scrap prices.

So I'll focus on describing the model for firms' expectations and the dynamic problem that the firms face. And, so, Zt is the demand state for the Asia-Euro market, and Zt-tilde is for the outside market. And Asian firms here consider an AR(1) process for the demand in the Asia-Euro market and the outside market, so this is the how demand states evolve over time.

82

And, so, the assumption of this learning model that I consider is that the parameters in this -- in these AR(1) processes are unknown to the agents. So agents update their beliefs by reestimating these parameters in every period using the demand realizations up to that date. Okay? So -- and, again, I'm going to consider -- or I'm going to allow agents to put heavier weights on more recent observations, and so consider various weighting of the past observations.

So in the figure that you're looking at, the case with the flat line on the very top is the case where agents put this equal weights on all observations and other cases where the weights are falling dramatically with the age of the observations. Okay? And this is, again, the case where agents are concerned about structural breaks and unknown dates.

So firms decide whether to invest and also whether to scrap their ships in each period. The state that they -- the pair of relevant variables are their current capacity, their order of book, that's how much they're waiting to get built, and the sum of everyone's capacity in the market, as well as some -- the industry order book. Also, there are two demand states.

83

In the static product market competition stage, firms choose how much to charter, that's how much to lease from third-party companies, and they choose how much to deploy in different markets. And they face constant elasticity demands for shipping services.

So here is our implement -- the model of firm beliefs. Now, for each of these weighting parameters, Lambda-t is the parameter that governs how much firms discount older observations. So for each of these parameters, I estimate the parameters in the AO(1) model using demand realizations up to that point. So this would correspond just to fitting like least squares on the growing sample and weighting -- weighted least squares if you have a case where agents are discounting all their observations.

And, so, what you're looking at is the estimates, the beliefs under learning model for the Asia-Euro market for one particular value of Lambda. And, so, the figure on the very left side shows the volatility estimate. So you can see that it jumps dramatically around 2009 -- 2008 or '09. And, also, the persistent parameter in the AR(1) model, which is the Rho-1, tends to fall steadily after 2007. Okay? And how much this volatility measure jumps or this

84

just persistent parameter changes is going to depend on Lambda, which is how much they discount their older observations and also the model of firm beliefs, and this variation will help me identify the model. Okay?

So the estimation proceeds in different steps, but I'm just going to focus on the last step, where I basically estimate the dynamic parameters and the parameter in the model of firm beliefs. The other steps are quite standard, but I just want to point out one thing, which is that I estimate the investment costs and scrap values outside of my dynamic model using this shipbuilding cost data and then do the last step.

So I use the method of simulated moments to estimate this model, matching the moments in the data including the average investment in the period before 2008 and after 2008. And the total capacity in the industry, total capacity in the order book, and the correlation between demand and investment and the volatility investment.

So the main result is that I find that the adaptive learning model where Lambda-t is equal to .02 fits the data the best. It's sort of interesting to think about this, so there are a couple of other papers that try to estimate this model in the macro literature using either survey data on some microdata on expectations, and it seems that the value that I estimated is quite close to their estimates.

One other parameter that I estimate in this step is a fixed cost, though this is fixed costs of holding onto capital. That does not vary with how full the ships are. So it would be maintenance costs, port charges, or, like, labor -- basic labor costs. And it's substantial. So it's going to be about 36 percent of period profits.

So here are just -- it's just the model but in terms of the yearly investment. I'm just aggregating at the year level, and the solid line is data, and then the line with circles is the model predictions. And as you can see, it does a pretty good job at predicting the boom in 2007 and then also the bust afterwards.

So I just want to briefly talk about the alternative models that I consider. So the full information benchmark is the one where parameters in this AO(1) model are known to the agents. And, so, here as a researcher, we would estimate this model using the maximum data available to us, so the full sample of data. And they endow those beliefs to the agents. I also consider Bayesian learning model and

also some more flexible specification of the full information model.

So here are the model fits under alternative models. Again, the solid line is the data, and the line with the circles are model predictions. And I just want to draw your attention to the figure on the left side. So that's the full information benchmarks case. As you can see, it does a really poor job at predicting the correct timing and quantity of investment. So as you can see, the firms are actually investing less during 2006 and '07, the investment boom period.

And, so, why -- what's driving this? Basically there were two forces that are going on. When demand increases. This has two effects. One is that, of course, the returns in investment gets higher and firms want to invest more. But at the same time, demand for ships increases, this raises the price of ships, and that's going to decrease investment.

And, so, in the full information case, this negative effect dominates and actually the correlation is negative between the investment and demand. In the learning case, when demand is good, agents also become more collectively optimistic so that this positive effect is going to dominate. As you can see, Bayesian

learning model does a better job, and also the flexible model of GARCH, but not as well as the baseline learning model that I showed you.

Okay, so, the remaining time, I want to talk about counterfactuals that I think about. So the first one is about competition. And, so, the question that I have in mind is whether strategic incentives increase the level and the volatility of investment and what happens if we increase consolidation in this industry. And, so, why do I care about this? First of all, there is quite big theory literature on how strategic incentives such as business stealing effect or preemption effect can also lead to overcapacity. And, also, in this industry, there has been a trend towards consolidation. So there's all kinds of proposed mergers and alliances that are happening.

In the model, there are at least two sources of strategic incentives. So, first of all, as a firm, as I deploy more capacity, that's going to increase my own market share, but it's going to have a negative effect on my rival's profits and market share. Okay, so that's going to lead to the business stealing effect.

But, also, when I increase my order, that's going to increase the aggregate order book, which is

going to raise the shipbuilding prices. So this is going to lead to the preemption effect where when demand is good, I want to be the one that's first to invest. Okay?

So the two things that I consider is that monopolization, which gets rid of strategic interaction between all firms, so it's a multi-plant monopoly where the market shares of the firms are fixed, but they, you know, make coordinated investment decisions. So that's the line in the bottom. And then the intermediate line is a merger case where I allow the merger between the top two firms. And then the top line is the baseline learning case.

And, so, what I find is that both monopolization and a merger decreases the level and the volatility of investment. So in the monopoly case, something like 34 percent and the volatility goes down by 21 percent.

So in terms of the welfare, what does it imply? It leads to a huge gain in producer surplus -- for the producers and some consumer surplus loss. So I just want to point out that the consumer surplus is incomplete, so it's only with respect to one big market, which is about 30 to 40 percent market share in this industry. But nonetheless, if you look at the merger case, it's likely the case that the producer surplus gain is going to dominate the consumer surplus loss because they're going to be not only reducing investment, but these merged firms will know when to invest. So they're going to try to invest more efficiently when price is lower, okay?

And, so, the last thing that I consider here is whether the modeling choice would matter when I do this type of policy valuation. So I do this merger exercise under the learning model and under the full information model and find that the learning model predicts a much higher change in the investment -- both the changes in the investment rate and also the welfare.

So the rough intuition is the following: When there is high demand, that's when there is high incentives to steal business or preempt your rivals, but under learning, agents are also becoming more optimistic during this period. So learning reinforces this preemption in business stealing effects and so intensifies the strategic incentives. And, so, it's going to predict a larger welfare gain from this merger. In other words, the full information model underestimates welfare gains from this particular merger.

So the last set of counterfactuals that I want to talk about today is with respect to demand volatility. And, so, I do this exercise under the learning model and full information model, and in both cases, increasing demand volatility is going to decrease investment slightly. And, so, this is actually consistent with some of the previous studies, like Bloom in 2009 and Collard-Wexler in 2013 that shows that increased volatility reduces investment.

However, if you look at the volatility of investment, you will see that there's also -- it's also the case that when demand volatility goes up, the investment volatility goes up. So there are two reasons why this is happening. First of all, as demand volatility goes up, the price of input, the shipbuilding price is also becoming more volatile, so that's going to lead to more volatility in investment, but also under learning, there is a second channel where higher demand volatility leads to more drastic and more frequent revisions and beliefs, and this is going to lead to more volatile investment.

Okay. So I just want to conclude by saying that, okay, this paper analyzes boom and bust cycles of investment under demand uncertainty. It builds an estimate of the dynamic oligopoly model with

uncertainty about the demand process itself. I show that a learning model can help us understand this type of firm behavior where they're investing a lot when investment is expensive. And I also show that strategic incentives increase the level and the volatility of investment learning sort of intensifies these forces.

And, lastly, I show that the modeling choice for firms' expectations has policy implication. Thank you.

(Applause.)

MR. RAVAL: So Allan is doing double duty today, and he's also going to discuss this paper.

MR. COLLARD-WEXLER: Okay. I'm not thinking what if I had presented after the discussion. Anyways...

So I want to start off by telling you why you should care about container shipping. So there's this great book by Marc Levinson called The Box, which, you know, read that, and, like, don't read my paper, be like read that book first because it's amazing. The kind of transformation of international trading relationship because of container shipping is enormous, and it's beautifully documented there.

And, you know, what I think is interesting

is most of you probably know Myrto Kalouptsidi's work on bulk shipping. And, you know, you might think, well, oh, this is another shipping, you know, paper. But bulk shipping and container shipping are really different, so my summary here is bulk shipping, like shipping coal or whatever, phosphates, from one place to another, that's like a taxi. You know, you just call one up, it's perfectly competitive. And container shipping is like airlines, so there's regular routes, and the issues of market power are first order, and there's been some cartel activity on -- in this world. It's quite concentrated.

So I think -- I mean, Myrto's even told me this, that, you know, container shipping of the two is the more interesting part of the global shipping kind of industry. And, so, I think this is why -- I think this is why from an antitrust perspective we'd really like to understand this market. And, you know, like other large commodity markets, it's had very large swings in total capacity, and people can tell you what China was expected to produce or not and what that did around the financial crisis.

So there are also this large amount of cyclicality in this industry. So let me just tell you, you know, what are the -- the components of this paper are, you know, a computational Ericson-Pakes dynamic oligopoly model. There's some kind of -- there's limited data here, and she's doing like a combination of estimation and calibration and kind of dynamic estimation to get different parameters of profits, investment costs, and so on.

And then the other twist is that rather than just using kind of simple Markov process for demand, she's using an adaptive -- adaptive expectations model, where the process gets updated when you see different realizations. So this is what this paper is combining in the model section.

So the first thing I want to push on is that it strikes me that what has to be done here is taking a lot of different price information and kind of reducing it down to kind of a single price on Asia-Europe in terms of that market, which now that I've told you that ship -- you know, container shipping is like airlines, you realize that there's going to be some heroic assumptions that go into that.

And, so, there's a lot of aggregation across, you know, this is a spot contract, this is a charter contract that needs to be done, and, you know, I think it works better as a "we really want to understand shipping, and this is going to be like the

market for taxis in New York," which is just a great example. You know, we do this because getting the numbers right is going to matter, rather than because the data is particularly great.

Volumes are easier to get information on. You see the ships. You see how loaded they are. And the demand system, basically, in my mind, I was thinking like Porter '83, you know, a CES demand and then some kind of increasing marginal cost specification.

And, so, the first question is just, you know, how much are we losing by this aggregation into a homogenous product to Asia-Europe. And this got me more worried because you've got this outside Asia-Europe market as well. And, so, if my first model was that, whatever, you could just put all these ships together, put all the demand together and that's like a homogenous good, then I don't understand why there's two markets that you're focusing on.

So these are heroic, but this is the only way we're going to get there. And, so, it just -- whatever the assumptions are, it would be nice to know what's the violence of the data that's going on here. But this is -- this is the first paper to attempt this.

I'll move on. There's some nice bits in terms of how the model is being estimated and solved. So, again, unlike Myrto's work, you know, there's companies like Maersk that are, like, whatever 20 or 30 percent of the global container ship volume. And, so, you know, the state of the market is going to be, you know, how big these different firms are and how many ships they have.

Now, the problem is if you're going to keep track of 10 firms' capital, which is how many ships they have, you're going to quickly run out of space in the computer to keep track of what everybody is doing. The state space is too big, and so there is some nice stuff on moment equilibria from Ifrach and Weintraub that's being used, you know, quite carefully in this paper.

I like this bit. My one piece is is that what's happening is in the paper I keep track of how many ships I have and then what's the total amount of capacity to everybody else has, so that's an approximation. And I think this industry -- you know, this kind of technique needs like an industry standard of how do we check robustness. You know, I could equally do one where I don't keep track of everything. You know, that's also a moment-based equilibria, but

maybe I don't like it.

And, so, being able to tell us why I should like using total capacity as the state of the rest of the market would just help evaluate, you know, is this working or not. And I just think we just need to get used to doing that when we're using these methods. Just we haven't used them that often so far.

All right, but again, this stuff let's you estimate a game for this industry with lots of -- lots of firms and concentration, and that's why she's doing it. And then I'll just say two things about the estimation. So one piece that she's doing is she's using the prices for scrapping ships and also the prices for ordering new ships to basically pin down entry and investment costs. And then the things that are being estimated are these variances of those scrapping costs and investment costs. And you can think about it that, you know, she knows the mean, but she wants to get the elasticity of, say, entry with respect to profits. So you need some variance around the scraps to get an elasticity.

So it's really the right thing to use the dynamic model to get elasticities rather than means, given the data is so good. One comment would be if these things are moving around year to year, is that

97

part of resolving -- you know, once I start thinking about endogenous input costs, you know, should that also be part of the model, given that they move around so much from year to year. But, again, using time-varying entry costs because in this industry that's a big deal, that -- I think that's a nice innovation here.

Let me move on to the learning process. So this seems -- this adaptive learning process has been around forever, so Jan Tinbergen's work has that from like the thirties. The caveat you should know about is -- and correct me if I'm wrong -- but firms don't have any awareness that they don't know things. It's just they had, like, one AR(1) process and then they updated, but they're not thinking about, well, maybe I don't know the AR(1) process. So it has some severe kind of constraints on, you know -- it's not like the uncertainty is part of the state in this model. Let me put it that way.

And the thing that this thing does that you might not be aware of is that you need, like, slow updating. And Bayesian learning models do very badly when macroeconomists have used these. You know, the shock happens, everybody gets it, and then it's over. Here, you get some kind of persistence, so there's --

98

this is not just a simple functional form. It also kind of yields things that are nice in terms of how quickly kind of new information percolates into the economy. And, so, I think that's what's good.

If there's any way other than using the full structural model to validate your estimates, that would be great, like resale value of ships or other -- there are some surveys, but not for shipping in this paper, but it -- at some point, there was a kind of a "how do I know this," other than the, you know, GMM criteria is a little bit higher for Lambda equals .002 versus .00, for instance. And, so, that -- I think that would help kind of shore up the evidence there.

So there's a number of different counterfactuals -- mergers, demand fluctuations, scrap subsidies -- so also related to different papers that have come before in the literature. I really think that, you know, the three -- the merger, you know -- the thing that I think this paper wants to do is distinguish what are the counterfactuals that I need the learning model for, and what are the counterfactuals that I kind of could have done without it.

And likewise, what other counterfactuals, what I could have used kind of a competitive model,

99

whereas what other the counterfactuals we're having in this whole oligopoly interaction means something, because that's -- that's really what this paper is combining. And, you know, I have to say of all of them, you know, it's really the does cyclicality of the industry change when you have -- when you have, say, investment, when you have a merger, given this learning story.

You know, that's the one that I think really combines the two pieces very nicely. And just given the amount of work put into getting these two components -- dynamic oligopoly and learning -- into one paper, you should kind of focus on the thing that -- at least I would focus on the thing that kind of brings them together.

One of the -- you know, some questions. So, like, you're estimating the learning model to get a sense of the -- and then using the estimates from the learning model, you're looking at the prediction is a rational model. It kind of struck me that if I'm worried about any type of misspecification of using the wrong parameters in the rational model, like, what would that do. So I think one suggestion here is just if you estimated the parameters with the no learning model, you know, how would those predictions work.

100

The other piece, and I don't know if there's much you can do, but it would be nice to get some sense of, you know, does parameter uncertainty in the sense of statistical significance of these things, does that affect these comparisons that you're doing. I just can't tell if the numbers are largely different or small. And, so, just getting an idea of, you know, if you drew from the distribution of parameters in some way, the dynamic parameters are hard to draw from because there aren't standard errors on those, but if you drew from all the other parameters here, would these effects be robust. And I think that would help kind of emphasize what's going on.

There's a lot to like. We want to know about container ships, and this paper does a good job at getting a first pass at what we can learn from this market. There are some nice things about the modeling of the industry, large state spaces, time-varying costs of ships. And there's also this kind of learning model which just says, you know, we can -- we can be more flexible about the other components of these dynamic oligopoly models and, you know, accounting for these kind of large swings in investment might be one way we can make our dynamic oligopoly models kind of match data better.

101

Okay, and that's it.
MR. RAVAL: All right, we have time for one last question.
MS. JEON: If there is no question, I can just address, like, some of the comments that Allan gave. Thank you so much again. I just want to say in a very limited sense, I did look at robustness of, like, changing the moments that firms care about, so I try to put in, like, the biggest firm states into the state space. And it didn't seem to change much.
And then, oh, for the full information counterfactuals, I reoptimize everything so that I found the parameters for the -- that works for the full information case, but otherwise, really great comments. Thank you.
(Applause.)
MR. ROSENBAUM: All right, thank you all. We'll take a 20-minute break and come back for Steve's keynote at 11:40.
(Recess.)

102

KEYNOTE ADDRESS:
"MARKET STRUCTURE AND COMPETITION, REDUX"
MR. ROSENBAUM: We'll get started.
Okay. Steven T. Berry is the James Burrows Moffatt Professor of Economics at the Yale School of Management, a research associate at the NBR, and a fellow of the American Academy of Arts and Sciences. He specializes in econometrics and industrial organization and is a fellow of the Econometric Society and a winner of the Frisch Medal.
Berry has previously served as the chair of the economics department and director of Division of Social Sciences at Yale and received his undergrad degree from Northwestern and his Ph.D. from the University of Wisconsin at Madison. Most significantly for me, he was my dissertation advisor.
Steve.
MR. BERRY: Significant for me, too.
Okay, so we've seen two nice keynote addresses, and they were, I think, the very best kind of keynote address, where someone gives us a concise, 30-minute summary of a body of their research. And it turns out I don't have a body of research right now I want to summarize in 30 minutes, so instead I'm going to give -- I'm going to give an old man speech, which

103

is along the lines of what's going on and what should we do about it in our field.
Very good. So when you talk to IO economists, there's something a little surprising, which is that of all the applied micro fields, in many cases, we're the one that often gets the least attention in the newspaper, that there's a lot of talk about, you know, taxation and public finance and, you know, labor issues and the minimum wage, and yet if you looked out right now at the debate in the paper, you'll see all kinds of things saying that the country is in the midst of a crisis of market power.
Joe Stiglitz has an article just entitled "America has a Monopoly Problem." The Council of Economic Advisors, you know, put out a report on this about, you know, the problem with markups. And, you know, you see kind of debates about hipster antitrust and IO economists have noticed this, and Carl Shapiro, for example, is giving, I think, a nice set of speeches where he says what should our policy response to this be.
And I want to talk about something more boring but much more dear to my heart, which is what should be the response of empirical IO economists, how should we think about the questions which are being

104

raised, and is it true that since we are, after all, the world's experts in markups that we have an answer for the questions that are being raised so prominently in the press and the policy debate and by our colleagues who are outside of industrial organization.
And I think the answer is so far we have not answered this at all, as near as I can tell, within empirical IO. And it's a little surprising. I've said to some people, it's a little bit like, you know, someone finds out their wife has cancer and runs to the biochemist next-door and says, you know, can you tell me how to treat the cancer, and the biochemist says, well, you know, I don't know, but there's this protein I'm investigating and maybe 30 years from now I'll tell you something about whether there's a treatment there.
So people have come out and said that maybe there is this aggregate problem in the economy that markups are very high, that we -- as Stiglitz says, we have a monopoly problem, that we have a market power problem and important enough for the Council of Economic Advisors to issue reports about it, and we have almost nothing to say about it. When I ask my colleagues in a room like this, you know, are markets -- are markups going up in general? Right, in

general, are markups higher in the economy -- we're the markup experts, right, that's our field, that's what we study. We study, like, pricing and markups and competition. Are markups going up in general in the economy, and they say the same thing that I do. They have exactly the same answer I do, which is how would I know that.

So it turns out, though, people are interested in this question, and they're doing a ton of research on it, and they're publishing a lot of papers. And these papers get hundreds of citations, and these papers are almost exclusively by non-industrial organization economists. They're macro people, they're trade people, they're labor people who spin a big theory about competition, and they collect some aggregate data and they find some correlations and some regressions. And they will give an answer to the policy people.

And the question, I think, is whether -- kind of how do we respond to that. Do we just say, well, it just turns out that part's macro and we'll just stay silent? Or is there some kind of response that we should have?

So one of the things I want to point out is that a lot of these papers by non-IO economists

recreate various aspects of the old and supposedly discredited structure-conduct-performance paradigm, which I'm actually old enough to have taught. I don't know if anyone else in the room took a course both from Mike Scherer and from Len Weiss. If -- very good, yeah. I was going to say, if Mike was here, I got at least 50 percent.

And, so, a few of us remember that and remember actually that it had some strengths, even though it got -- even though it got killed off. You know, so, you know, I want to talk about that, too, sort of how should we think about the use of techniques that would seem very familiar to empirical industrial organization economists of the 1970s, and here we're using the second decade of the 21st Century, and these are the answers which are -- these are the methods and the answers which are being presented to policymakers.

So, you know, it's almost a little worse than I thought. I actually kind of like this Autor paper on superstar firms, although, you know, from the IO perspective it has some crazy elements to it, but, you know, one thing is I just looked through for the cites to empirical IO. So, you know, there's Demsetz 73, Industry Structure, "Market Rivalry and Public

Policy." And there's a Schmalensee paper from the '80s that's another structure-conduct-performance paper. And other than that, I mean, I didn't -- maybe I missed somebody, but I don't think it cites any empirical work by current members of the NPR program except in reference to the estimation of productivity. It's a paper about competition and markups. And their claim is we have nothing to say, or at least nothing that's worth citing when they write the paper. And it's a pretty well-known paper.

Okay, so, part of what I want to do is I want to think of how should we think about, well, you know, this new structure-conduct-performance literature, which is being kind of reinvented by people in other fields. You know, so what could we do with it? We could ignore it. That would be one thing to do. We could pretend it's not happening and just say our -- you know, the way we treat a lot of macro, like, wow, interesting things happen in macro, that's crazy, okay. Wow.

Maybe I'll collect some more data. We could critique it. Okay? We could say -- we could remind them why we thought this was bad or at least try to say what the pitfalls are of doing it. And in some sense, maybe try to take the literature back to where

it was in the late 1980s where the structure-conduct-performance people were trying to improve their regression before they got buried under a tidal wave of a game theory and empirical IO disappeared for five or ten years, only to come back in a different form.

We could talk about improving it. Are there aspects of it or some parts that are better than others? Maybe we'd actually like to be a little more positive than critiquing it and say, well, maybe you should do this rather than that, or here are the things that are particularly bad, but here are the other things that make sense. We could try to improve it.

And/or we could propose alternatives. We could say actually within modern empirical IO we would actually -- we confess these are good questions, and here's how we think we should go about it. Okay.

So I'm going to take all of these bullet points seriously except for the ignore. I just think that's a mistake. I think, you know, we should care about questions about markups in the economy, in the American economy, in the world economy, and so I want to think a little bit about critiquing it. I want to think a little bit about improving it, and I want to propose some extremely tentative alternatives and maybe just do that as a way of trying to get some conversation going. Okay.

Okay, so what was structure-conduct-performance? As I say, some of us are old enough to remember it, and then some of you are young enough that they still taught it to you in a graduate class, and some of you may be young enough it never came up because it's why should you study economic history in a second-year graduate class.

I would say the broad question here is very much the same question that a lot of the papers today are answering, and it was asked for the same first-order important reason, which is people wanted to know what is the effect of market structure, often called concentration, on various outcomes, which were most often prices or products or profits but other aspects of conduct in the performance of the industry. And I'm going to say causal effect, which people in the '70s would not have said, or the '80s would not have said, but I think that's what they meant.

They meant it in the same sense that Josh Angrist means causal effect, right, that there's a Y variable like price or markup and there's an X variable which is concentration, and I want to know the causal effect of X on Y. I think that was very

much what that literature was about, right? And then you could ask, you know, what are the problems with that.

But it seems like a decent question, which is why it dominated empirical IO for a few decades and generated -- again, Mike and I had Len Weiss. I think he counted at one point something like 2,000 published structure-conduct-performance papers in the literature.

What was the typical method? The typical method was cross-industry, usually OLS regression of -- I think I actually reversed my sentence there -- of accounting measures of markups like the Lerner index or profits and other market outcomes on the Herfindahl index, which would be treated as the market concentration measure most commonly. And we can come back to why that was done.

So a classic regression would be an accounting measure Lerner index on the Herfindahl index. And you want to know is the coefficient positive, and if that is, that means that concentration raises markups. That's the causal effect of concentration on markups. And you could have a bunch of controls. Again, maybe you don't want to use Lerner -- accounting -- an accounting Lerner

index, which was typically revenue minus variable cost over revenue. Maybe you want to use a direct measure of profits. You know, maybe occasionally we had price. You could also think of other market outcomes, you know, and put them on the left-hand side of your concentration regression.

Now, at the time, it was controversial, even though there were 2,000 -- even though there were 2,000 published papers. It was, nonetheless, a controversial thing. And a lot of the controversy at the time focused on what was called the Chicago critique. And, you know, there are various ways of thinking of the Chicago critique, but a lot of it really had to do with the theoretical endogeneity of market shares and that if you think of the Herfindahl index, which is, in fact, just a transformation of the market shares within the industry, so somehow the market shares are leading to concentration, right? And Chicago liked to emphasize reverse causality, that if you get a big firm, you know, a Cournot model or in lots of models, that would be a low-cost firm. Low cost leads to high shares for that particular firm.

So take an industry that's relatively deconcentrated, have one of the firms invent a new technology, which makes them much more efficient, it

drops their marginal cost, the Herfindahl -- that firm will gain market share and now you'll have an asymmetric industry with an efficient firm and a bunch of inefficient firms. That share will go up and the Herfindahl index will rise. Right, the Herfindahl index will rise. So they said really all of this concentration is a result of reverse causation, is really about endogenous market shares. And sometimes even though they had said they were endogenous, they would do things like regress a firm-level Learner index to the market share, right, and say, well, gee, firms with big market shares have low markups.

Now, again, if it's theoretically endogenous, there's a question of why he just ran OLS, but that's the kind of thing people would do. What were some other critiques? And another one that's come back? You know, accounting data are terrible in many -- in many ways. There's a ton of mismeasurement, capital is very -- extremely difficult to measure. Everything's aggregated. You don't have detailed product measures. You often have no price variable at all. You're only seeing revenue, which is why you get some kind of accounting profit or accounting margin that you're using because you don't know how to have a cross-industry measure of price.

113

So there was a lot of the critique of structure-conduct-performance would be about the problems of accounting data. Now, the advantage of accounting data -- and what I'll come back to is that it exists across industries and across -- so if you want to do something cross-industry, it's very difficult to avoid accounting data because there just aren't consistent sources of data that, you know, have price and things that modern empirical people use, right?

Another critique is that there was really no single cross-industry theory of markets, and I think that's a hallmark of a lot of empirical IO people. We don't really think there is "the" theory of "the" market, right? We think that you have to match the theory to the market, that different things happen in different places and sometimes product differentiation is important and sometimes it's not.

And sometimes there's collusion and sometimes there's not. And sometimes capacities are really relevant, and sometimes they're not. And sometimes the dynamics of the market are important, and sometimes they're not. We think there's a just different theory for every market, and so how do you possibly run a cross-industry regression when you

114

don't even have a single theory that you think runs across markets, right?

And, of course, implicitly in structure-conduct-performance, there was something like a Cournot model always running behind it, which I'll come back to. The question is how bad is that. Okay, so Mike and I had courses with Len Weiss. Len, by the end of his career, was trying to save structure-conduct-performance, said, okay, you know, Chicago says shares are -- and Herfindahl's are theoretically endogenous, you know, let's treat that seriously as econometric endogeneity and, you know, now maybe we're really, really, really in Josh Angrist's world where we have a Y and an X and a Z, right? So we're looking for the causal effect of X on Y, but X is endogenous, and so we have an instrument Z, and could you possibly do this as a Y-X-Z model within instrument Z. And what would the instrument be, right?

And, so, I want to come back to that a little bit. Is that a possible solution that we have, you know, kind of instruments or how often is it a solution, at least to this -- at least to the Chicago school -- at least to the Chicago school problem.

Okay, so what was the -- what was the -- what happened? As I say, well, game theory came in,

115

and then Bresnahan proposes the acronym that fell out of use, the New Empirical IO, partly because it wasn't new, partially because NEIO is not a great thing to say out loud anyway. And what he suggested was that we have single-industry studies. Why? Because you could get much more carefully measured data.

You could actually maybe get price within an industry. You could get price separated from quantity within an industry. You could start to get product characteristics. Occasionally you might get cost data. You would know what theory to tie to this market, one would hope, and you could start putting it in an oligopoly context, which was closer to classic supply-and-demand analysis in the sense that your analysis of endogeneity and identification and instruments could be based on, you know, ideas that demand shifters are excluded from the cost function or cost shifters are excluded from the demand function and so forth.

And you could go back to a much clearer kind of classic supply-and-demand style notion of instrumental variables and endogeneity and identification. So only for the purpose of these slides and nowhere else I'm going to jokingly call this the dominant empirical IO algorithm. I don't

116

think we should adopt that anymore than NEIO, actually.

But this really is the dominant algorithm, which -- within empirical IO in a broad sense that we're single -- we're crafting studies of single industries with the theory guided toward that industry with an industry-specific data collection and identification tied to the cost and demand shifters and equilibrium assumptions of that industry, and we're handcrafting all of these little individual industry studies, right, in a way that we think is, you know, other people say we're making too many assumptions, but in a way that we think is pretty careful, say, compared to structure-conduct-performance.

So my colleague, Bill Nordhaus, pulled me aside one day, and he said, you know, the thing about you guys is it's like -- it's like house-to-house combat for you guys. You know, it's just like -- there's this big battle and you just took like a house. He said, you know, macro is -- in a good way, it's like carpet bombing. We just solve all the problem at once. I'm like, okay, is that good, carpet bombing? But, okay. It kind of is the way I think about macro, but okay.

You know, but here's the critique, and I think we should take it seriously, which is that macro studies the economy and we're interested in a particular part of learning about marketing for yogurt. Right? The Journal of Economic Perspectives, once asked me to write an article on what has IO learned about markets in general.

They're like, oh, you can't write that, huh?

Now, you know, okay, the interesting thing is while the macro economists are moving in behind us with structure-conduct-performance, I actually think we've done a good job over time of going into markets that are actually pretty important and incorporate a bunch of the economy, like health and education and environment and in addition to the broad studies of antitrust that ought to be directly relevant to these questions that everybody is asking right now.

So I don't want to -- you know, I actually think we've done a lot of important of policy-relevant stuff, and in some sense, as we have been sort of colonizing big areas of what used to be public economics, precisely because we can do supply and demand equilibrium studies that used to be the theoretical hallmark of public economics, they -- public economics, meanwhile, has adopted this kind of

Y-X-Z strategy, where you're not doing equilibrium studies, but you're looking for causal effects. And I think it's been a big success.

But, you know, despite that, I think actually pretty big success and that kind of policy success that has made IO I think increasingly policy-relevant in many parts of micro, very few people, and actually I say Allan's a bit of an exception here, have been doing things that are sort of geared toward the macro aggregate conversation about markups and competition in the economy as a whole, and I think this is why people have gone back to structure-conduct-performance because that's -- that literature was an attempt to answer that question and maybe is the first thing you would do if you weren't steeped in modern IO and wanted to answer that question. It's just not -- not as insane as, you know, I'd sort of like it to be.

Okay, so, you know, I didn't do a full literature review, right? You know, I think -- on an earlier slide, I think I had that, you know, you just look up that Council of Economic Advisors report and, you know, just alphabetically you get a reference, you get another reference. Both of those are regressions on Herfindahl, right? Modified Herfindahl, innovation

on Herfindahl, something on the Herfindahl, something in concentration, right? Just stop -- at the moment we -- most of we in this room would not do any more.

I guess the one -- I'm sorry, the one on the earlier slide was that, you know, that -- the Demsetz thing, that, you know, they're just not citing us, and they are citing the structure-conduct-performance literature or this modern structure-conduct-performance literature that is stuff on Herfindahl's, stuff on H.

I really thought that that -- that that highly cited paper on innovation under Herfindahl, I was really pretty sure that was in Mike Scherer's 1980 textbook, but it's a fat book and I couldn't find it. Maybe he just sketched it on the board, but I don't know.

Okay, so, you know, not all of these are all -- not all of these sort of new structure-conduct-performance papers have all the features of structure-conduct-performance, but, you know, they have some of the features and/or all of them, which is you're using cross-industry data rather than, you know, our strategy of single-market data and/or they're using accounting data so that they can run cross-industry rather than kind of carefully crafted prices and other

things. They're throwing concentration in. They're just trying to get direct measures of markups, maybe. And sometimes they're treating market structure as exogenous, as just this exogenous thing that's out there, even though they're using H; and/or sometimes trying to use ad hoc instruments. In other words, they're sort of doing the Y-X-Z thing of recalling a variable that was not in X and sort of conjuring it out late in the paper and saying, oh, it's an excluded instrument and you're not quite sure why except they forgot to put it in earlier, and so it wasn't in X.

Okay. So, okay, here's the critique. I still think that as well as intentioned as it is that straight-up causal effect structure-conduct-performance is still pretty hopeless, right? I don't think you can directly say this. You know, so let's say you had a really good price or a markup, right, and you're going to regress it on H to get the casual effect, but what is that thing? It's not a demand curve. It's not a cost function. What is it?

I think it has to be the first-order condition from an oligopoly model. What else creates the relationship between price and fundamentals in an oligopoly? It has to be the first-order condition from an oligopoly model. Now, just thinking about them in general, without the structure of a particular model, first-order conditions include the effect of demand and supply because there are markups, and markups depend on demand, so demand is in there. And prices also depend on cost, so cost is in there.

So what is excluded from that relationship that could possibly be an instrument? Right? People -- I've been at talks; people say, oh, don't worry, I instrumented for H. This just happened at the Searle conference. Don't worry, I instrumented for H. I understand it's endogenous; I instrumented for it.

What variable is excluded from the price Herfindahl equation? Right? What's excluded from that? Okay, if it's a first-order condition, the pure measures of fixed cost, maybe, but I've never seen people do that. But fixed cost is excluded from pricing decisions. We teach the undergraduates that.

Something about exogenous merger policy, I guess maybe? I haven't really seen that done well. And part of that is that -- you've got to get a lot of variation cross-sectionally from these things. There's not that many time periods, and, you know, even if you thought the merge policy changed in year X, well, that's like one change in the variable. How much cross-section -- do you have real cross-sectional

variation in that? Right, if we had regional, I don't know, maybe if we had regional DOJs or something, policy instruments.

And just more fundamentally, there just isn't a direct model that has an effect of H in it. If you look at a sensible first-order -- and there's a second, if you just remind everybody, everybody knows this, if you just look at a model that has H in it because you did something that generated H, there is no effect of H in that model. H is just a joint endogenous outcome. As Chicago said, I don't know, did the demand elasticity go up, or did a marginal cost go down? Right?

That will affect price, and it will affect H, but not via the effect on H. Right? They're both just endogenous outcomes. Right? This is a lot like saying, you know, imagine another Y-X-Z paper, you know, I want to find out the causal effect on price of quantity, right? All the quantity -- and there are many theories of why quantity affects price. All the quantity's endogenous. What instrument should I use in the pricing equation? Right? Now, everybody knows that's a huge mistake. There's no such thing as the pricing equation, and there's one thing called a demand function, and there's another thing called a

supply function. There is no such thing called the pricing equation. And I'm going to argue there is no such thing as the H equation.

Right? So I really still think this is fundamentally a problem with kind of the regression on H. Okay, so let me hammer it home, okay? So, okay, so the only way you can get H in a model that I know of, and everybody knows this, is Cowling and Waterson in '76 or something, is via the Cournot model, right? So Hugo had the Lerner index, price minus marginal cost. The J -- M is market, J is firm. I should have had a J index on that marginal cost there, I see. I'm just multiplying his equation through by price so I get an inverse semi-elasticity instead of an elasticity.

Let me give a little econometric structure to marginal cost. Really that beta is not a coefficient that ought to be varying. It ought to be varying with market quantity. It ought to be varying with demand shifters and stuff, but, you know, given that it's constant if you're going to do structure-conduct-performance or something. It should still vary across every market, by the way. There' s no reason for it to be constant across markets, and I would get like a -- I would get like a Chicago

regression a little bit here, which is price on market structure, right?

But, again, is that the causal effect of market structure on price? No, it's just putting two endogenous variables in the same first-order condition. Right? Furthermore, demand stuff should enter beta, really. All the cost stuff is already in the equation. There's a huge endogeneity problem, and the Cournot model -- the cost shock is determining the share. Indeed, it's the only determinant of the within-market variation and share as marginal cost. Right, it's just one-to-one within market.

And everything else is demand, right? So, you know, this is Bresnahan's point. And by the way, you know, share is really just quantity divided by industry share. And if marginal cost slopes up, quantity also enters marginal cost. This is Bresnahan's point. Quantity is in there twice. Give me one instrument. Give me as many instruments for q as you want. How do I distinguish the effect via demand and the effect via quantity? That's Bresnahan's point. You can't distinguish these things. Bresnahan's point was, well, fix beta, and I can tell you -- and I can tell you the effect of q on -- fix the model, fix beta, and I'll tell you the

125

effect of q on marginal cost.

Without saying this a Cournot model, without saying that I have an estimate of demand to learn something about beta, and without specifying marginal cost, this equation is hopeless. There's no causal effect of share. Even the Chicago guys are wrong. There's no causal effect of share here. So what are you supposed to do? You're supposed to average that equation, and then you get -- you get concentration comes out.

Right? So if you take a simple market average, the average share is always one over N, right? So then I get an equation that relates price, I've got the same semi-inverse -- inverse semi-elasticity to one over N. I've got the average cost shifter now, and I've got the average output, right, to tell me to learn about economies are just economies of scale, and I've got the average cost shock, right?

So, now, the SCP folks hated N. And why do they hate N? Because every industry they looked at there was some gigantic tale of tiny, little firms and how do you count N. On the other hand, I got to say that N is responding at least at lower frequency to cost shocks probably, right? And I can almost imagine some instruments for N if I could measure it and some

126

correlates of W-bar. I object to this less than the Herfindahl one, but it's really got all the same problems, right?

And, again, N shows up twice. It's affecting economies or diseconomies of scale, and it's also having this competitive effect via the demand side moving down the demand curve in the Cournot model.

So I don't know. The next one is worse, though, which is the classic one, which is the Cowling and Waterson one, where you take a share-weighted average, and the share-weighted share is the Herfindahl index. I get the share-weighted cost shifter, and, you know, effectively the share-weighted quantity or you can think of that as Q times H, classic thing. And then I have the share-weighted cost shock.

But now that cost shock has shares in it, right? So now when one of the individual cost shocks goes up, mechanic -- when the share of a low-cost firm goes up, the weight you put on its shift -- on its weight goes up, which actually changes the shock. You get this mechanical relationship now between nu and everything else. W-tilde there, the share-weighted cost shifter is endogenous now. It's got share in it,

127

right?

The share-weighted first-order condition, right, is a tough thing to instrument because it's a function of shares everywhere. And this is the only way that anyone knows how to get H in a regression that has price in it. Okay.

It does not get easier with product differentiation or different models of competition or anything else, and this is why we gave up on it. And we were right. Okay.

Okay, one thing. If you want to regress price on concentration and just tell me it's the descriptive regression, I have to accept that because you have described the data to me. Right? So I would actually rather see the OLS regression of price or markup on Herfindahl and just say, look, it's a correlation, it's not a causal effect, and I'm just describing my data set to you in this particular way, and we can talk about what -- you know, if there's some model or something.

So I would rather you not instrument, right, and just give up on the idea that it's the causal effect. It's a descriptive regression, and it's a fine thing to do. And, actually, you know, Autor's paper comes kind of close to that, to tell you the

128

truth. They're thinking of a hidden third factor that's kind of moving both.

On the other hand, if there's a hidden third factor, maybe we should just look at the reduced form. Why not look at the reduced form effect of that third variable on price or markup and on concentration? Right? Why are we sort of going indirectly through these two endogenous variables? Maybe it's all we got. Maybe we don't see the third factor. But that's the only excuse I can think of.

Okay, I'm going to have two possible non-structure-conduct approaches to the question -- again, the question I think we're being asked, which is are markups going up in general. Okay, so, one, this paper got a lot of attention, Jan De Loecker and Jan Eeckhout. Okay, so they're going to do something a little SCP, because they're going to say just up front, and I've seen Jan talk about it, he says, I'm going to use accounting data. He says, I'm going to use the worst accounting data there is, which is Compustat data, and I have for years told my students never to use the crappy accounting -- Compustat accounting data. But, you know, manufacturing isn't that big a part of the world anymore, and this is accounting data that will tell you the economy, not just manufacturing. You know? So maybe you got to compromise sometimes.

And they're going to go for the macro markup, cost over price or price over cost. And they're going to do it without any of the demand data and without an oligopoly equilibrium assumption, purely through cost minimization. Right, so two things, accounting is not very good. Jan's very aware of that, and he's very aware of the other thing, too, which is the Chicagoans called from 1975 to say that high markups might reflect low cost.

Okay, so, you know, most people know this math. You start off with the Lagrange multiplier for the pure cost minimization of a variable input that caused the -- the Lagrange multiplier is Lambda. We recall that Lambda in the cost minimization problem is equal to marginal cost. And we get that marginal cost equals the wage divided by the marginal productivity of labor, and we just rearrange that problem, we multiply everything by L, and we divide everything by revenue, and we rewrite the whole thing and we take it to the other side, and we get this Hall markup, which is the input elasticity of output of the variable input divided by the input revenue share is equal to price over marginal cost. That's just a fact about

cost minimization, by the way. And the input revenue share is actually probably not so badly measured in the accounting data, and so kind of the only problem is that we need to know the input elasticity of a -- of a variable input. So Bob Hall would have said, well, it's cost that returns to scale, it's Cobb-Douglas, so the input elasticity is the cost share of labor, so it's the cost share of labor divided by the input share of labor, so it's just cost over revenue. Right? And that's kind of the macro approach, is marginal cost over price or price over marginal cost.

So the key question here is then -- it flips to a nice supplied micro question, which is this is a technology adjusted input revenue share, and the question I think is are we really sufficiently allowing for heterogeneity in technology, right? Because the question is going to be are prices changing over industry over firm over time. And we're going to get the right industry -- answer to the degree that we have estimated this input elasticity not correctly on average but correctly over firm and industry and time.

Okay. Now, I mean, this is kind of nice, and Jan's point is that markup is a residual here, just as in many of our models where we don't see any

cost data, marginal cost is a residual of the first-order condition. So it's -- you know, it's kind of nice. It's not really a dual, but it's a similar kind of notion.

As I say, it's really -- just as our measures of marginal cost depend critically on us getting the own firm elasticity or the cross-elasticities, this depends really critically on getting -- getting the beta right, getting the input elasticity right. And against that, you have to put the advantages of cross-industry data. This seems like a good complement to me. That's not really going to answer people's questions in the end, as John has said, because people want to know whether it's price going up or costs going down.

Right, but it's a nice complement. I mean, it's -- and it uses accounting data. What can you say? For example, what do they find? They get a big increase of markups, which we should at least say, okay, that's a possibility, that's what they found. Big increase in markups beginning about 1980. High-market firms tend to be smaller, which goes against a ton of other theories and makes me worry that they're not getting the input elasticity right for small firms. It worries about that.

And it's mostly within industry. So that's interesting because it doesn't sound like we're failing to, you know, enforce the antitrust law someplace and doing it other places or things like that.

Okay, I've got a minute. Here's the other idea that I just want to point out. So another thing is we could do what we do, but we could compromise a little bit, which is can we do some studies that are bigger aggregates of the -- could we do some studies that are on bigger aggregates of the economy and ask this question using our best tools. We might occasionally have to bring some accounting data, or we might not be able to do our fanciest model because we're going to do it -- you know, we're going to have to assume some things are constant.

Maybe the theory is constant across a little bit bigger set of industries than we, you know, had thought of. We have our workhorse industries. I don't even think we've done it within our workhorse industries. What about airlines? What about automobiles? What about the healthcare sector? What about supermarkets? Are markups going up there? We could at least say that. It seems like we owe people that, actually, if you ask me. I told Marty -- or I told somebody yesterday I was giving this talk, and I said it was going to be about what we should do, and he said, no, you mean what young people should do.

So one example of this would be we have a student -- and actually I think he's showing up in DC next year at Georgetown, he's at Dartmouth now -- who did -- took the accounting data plus some geographic data from the census of wholesaling and kind of just did really standard IO on it. And he starts with some interesting facts, which sound a little bit like what people are saying in a way. The wholesale sector is growing, by the way. I was super surprised by that. I thought Walmart had disintermediated the wholesale sector. It turns out it's growing really a lot.

There are fewer but larger firms. It sounds like an increase in concentration. They have many domestic locations. They're offering an increasing variety of products. That's interesting. That's not like mergers or something. That's maybe a better output. They're often sourcing both domestically and internationally, and there may be some fixed costs to that. Accounting markups are growing, and IT spending is growing.

Okay, so what story is that? So what he did is he just took a set of really standard IO tools and

in some sense just explained this data through the lens of those tools, right? And there was no alternative theories. It was not like testing something, really, more like a decomposition. If you took the standard -- really long standard IO stuff and just passed this data through those tools, what came out of the other side, right? So he's got nested logit demand and price-setting Nash, and there's geographic competition, and there are some fixed costs of variety and foreign sourcing and some other stuff and some really crude free entry model.

He doesn't solve all the endogeneity problems because you don't see the detailed cost and demand shocks when you enter, but like a lot of people do that. He does consider the endogeneity of the pricing decisions, supply and demand, and so forth. Really, really standard, standard, standard, standard, standard stuff.

And what does he find? You know, okay, you're selling more with bigger markups, and the model's going to explain that through an increasing demand for wholesaling, particularly for firms that have a lot of variety, a lot of locations, and also foreign sources as well as domestic, okay?

Marginal costs are decreasing. Prices

aren't going up that much. And the entry model, therefore, suggests increasing fixed costs. Right? Together, demand is up and costs are down. The markups are increasing a lot. Firm size is increasing a lot because of the importance of the fixed costs. And the markups aren't competed away. It's consistent with this being an effect of IT driving up the importance of logistics, fixed investments and software that give you better geographic cost as you deliver your goods, fixed costs of opening operations in China, and it's an interesting story. I think I just said all those.

And I think it's a good question of, you know, how common is this, for example. I suggested it for airlines a long time ago, that networking lowers marginal cost, drives up demand, drives up demand for some reasons which might be good and some reasons which might be more like marketing and bad things, right, but they both lead to higher markups. Increased demand, lower costs, higher fixed costs. You get higher margins in variable profits. Fixed costs are naturally limiting the amount of entry. Right? That would explain higher markups.

Is it true for a lot of industries? Could we figure it out for a lot of industries? And the

last point I want to make is that a lot of the interest -- if you look at the Autor paper, which comes very close. He has a quote, unquote theory of superstar firms, which isn't well elaborated, but it's a little bit like what I just said. And he says, okay, the superstar firms are employing less labor.

A lot of the interest in this has to do with distribution, which we might think of as input demand, and we have a tendency to skip over that. So what are the implications of our even market-by-market competition models for input demand, which is the -- which is getting toward the distributional impact. Are the returns to labor and the use of labor changing relative to the returns to software and capital and so forth? Those are questions it seems like we could answer maybe in there.

(Applause.)

MR. ROSENBAUM: All right, thank you, Steve. We'll take one question and then can continue the conversation after the panel.

Ginger?

MS. JIN: Thank you. I really appreciate the keynote here. I just want to ask probably a simpler question than you're asking. Does market concentration go up over time? This is not sort of a

137

question of what's the effect of X on Y, just sort of what's the trend on X.

MR. BERRY: Yeah, so really I think there are some competing papers on this. And, again, if you went back to Scherer's text book, that would be like, you know, table one, concentration, concentration over time, concentration over industry, right? We kind of stopped doing this a while back. So Autor's paper claims yes. I think the two-Jan paper, De Loecker and Eeckhout, claims no. So there's a measurement issue there, I think.

MR. ROSENBAUM: Okay, now I'm going to turn it over to my colleague, Doug Smith, for our final panel on privacy and data security.

138

PANEL: PRIVACY AND DATA SECURITY

MR. SMITH: Maybe the panelists wanted to sit down.

So thanks, everybody, for staying for this last session. Hopefully it will be a lively one. You know, privacy and data security is not really an area where I think I need to elaborate on why -- a lot on why people are interested in it. It's kind of a big topic these days, but the FTC has a particular interest because, you know, it falls under our consumer protection mission. And, so, you know, we're really delighted to have four panelists today who can speak to both the state of economic literature and also talk about their own contributions to it.

So before we get started on that, though, I just want to do a quick plug for our PrivacyCon, which is happening February 28th of next year. This is a one-day conference where the focus is on new research. And the -- actually the submission date is two weeks from today, so if you're interested, I encourage you to look into it quickly and whether you're going to submit or not, you know, it might be an interesting conference to attend.

So with that, I think I'll just plunge right into introducing the panelists. So the way we're

139

going to start this off is each panelist is going to get a little bit of time to discuss their own research. And, so, I'll sort of do the presentations as we go along -- sorry, do the introductions as we go along.

So first we have Frank Nagle. Frank Nagle is an assistant professor in the Management and Organization Department of the Marshall School of Business at USC. He studies the economics of IT and digitization with a focus on the value of crowdsourcing and cybersecurity. His work utilizes large data sets derived from online social networks, financial markets, cyber attack data, and surveys of enterprise IT usage.

Prior to his academic career, Professor Nagle worked at a number of startups in the information security industry. In these roles, he conducted red team tests, responded to credit card and intellectual property breaches, and developed a two-week course that all FBI cyber agents must pass before entering the field.

So please talk to us about your work.

MR. NAGLE: Great. Thanks, Doug.

So, yeah, so my work looks at the value of goods that have no price, which in the digital economy

140

is increasingly a lot more goods. So that kind of breaks down into two buckets. One is crowdsourcing, and the other is security and privacy. On the crowdsourcing side, a lot of my work stemming from the dissertation studies the value of open-source software, so free software. How do we value this at the macro level? We've looked at how it's -- the fact that it has no price, has weird effects on calculating GDP. We've also looked at the more micro level, at the firm level, of how using open-source software can impact firm productivity in positive ways but only for some subset of firms.

And then we've kind of dug in a little bit more where open-source is a crowdsourced good and firms can actually contribute to it, although this seems kind of counterintuitive because you're paying your employees to write code that your competitors can actually use. But what we show is that the firms that contribute actually learn while they're doing this, and they end up getting more productive value out of using their open-source.

And, so, now we're starting to do some more things related to regulation, technology procurement at the federal level, to better understand the role of the Government in these types of things. And we're also thinking about -- there's some bills before Congress right now, and there's a push from the White House that's been going on for the past few years to increasingly use open-source and open-source governance mechanisms as a way to increase transparency within the software supply chain.

And, so, this leads naturally into kind of a better security, if we have a better sense of what's actually being used in our firms, in our organizations, in our federal agencies, then we can better actually secure it and invest in the right amount of defense against these things.

And, so, on the other side, in the security and privacy side, as Doug mentioned, that was really my background before going back into the academic world. And now we're looking at some large data sets, about a hundred million observations of various security events against the Fortune 500 companies. And we're using this to show a couple things. One is the importance of actually fixing the low-hanging fruit, so simple things like patching and closing ports and having good password policies. As it turns out, those actually matter. And there are still a lot of firms that are not fully investing in those kind of low-hanging fruit as they should be.

And the other thing we're looking at is competitive response. So one of the things we see is perhaps unsurprisingly, but we've -- you know, nobody's shown this before, is that when a company like Target gets hacked, Walmart and all the other big retailers start investing more heavily in security. And we're trying to kind of tease out as to whether this is something that's just awareness, so, wow, Walmart knows that they can be hacked, or is it some other kind of raising the bar and something that they actually advertise to their customers, you know, we have better security than Target, so you should come shop at us rather than at Target. And we're digging into that right now.

So that's kind of the high-level overview of the things I'm working on right now.

MR. SMITH: Thanks, Frank. So that was great.

So, Sasha Romanosky is our second speaker. He's a policy researcher at the Rand Corporation and a former cyber policy advisor at the Department of Defense. He researches topics on the economics of security and privacy, national security, applied microeconomics, and law and economics.

His research has examined questions such as

whether or not state-level data breached disclosure laws and reduced consumer identity theft, when and how firms are more likely to be sued when they suffer a breach, and when they're more likely to settle legal actions. He's also studied the cost of data breaches in order to understand whether corporate losses are really as severe as is commonly believed.

And most recently he has collected a data set of cyber insurance policies to examine how insurance carriers measure and price cyber risk. So Sasha.

MR. ROMANOSKY: Thanks. So it's been an interesting exercise to try and summarize my body of work. It's probably something we should all do every few years. But as I was -- so actually, earlier today, as I was sitting and listening, I was doing a bit of that. Hopefully no one will begrudge me for it. But I think I'll characterize it like this. I think I started out being very interested in understanding different kind of policy interventions that can be applied at a federal level, even at a state level, to try and incentivize firms and consumers to adopt better behavior.

So firms invest in security. They have many different reasons for doing so -- regularity, peer

pressure, shocks to the industry, say because of a data breach, and certain different kinds of regulatory interventions. And, so, the way I have tried to characterize that, or at least the way I was framing it in my mind was in terms of just very simply ex ante regulation.

We're going to apply compliance regulations to these firms to try and get them to at least reach a minimum standard versus ex ante liability. We're going to allow the accident to happen, the data breach, the security incident, and create a framework for injured parties, consumers, to bring actions to make themselves whole, so these data breach lawsuits.

In the middle somewhere is information disclosure. So an event has happened. It hasn't really caused any kind of demonstrable loss yet, so we're going to inform people. We're going to empower these consumers. And that's where these data breach laws really fit in. And, so, I guess I've tried over the years to try and understand those different components to understand whether or not firms are really incentivized to do the right thing and are they actually doing that, how are consumers reacting to all of that, and are we better off by any kind of -- any measurable factor.

And, so, as you heard, I've looked at the effect of these breach disclosure laws on consumer identity theft, which led me into looking at the litigation. It was always the story that plaintiffs would bring these actions against firms -- class actions or just individually -- but they would always fail. And, so, I tried to understand, well, do these lawsuits actually fail? Are there any kinds of settlements? And what are the characteristics of a breach that lead to litigation? What are the characteristics of the lawsuit that lead to any kind of settlement?

And that was quite interesting. And that took me into the story of the cost of data breaches. If we think cyber is really a big deal, if we think these security incidents are really a big deal, like we always hear about, is that actually true? And, so, I was able to collect the data set to try and understand what these costs are. And what shook out of that was the notion that, well, maybe they're not quite as intense as we all think. From the data that I looked at, they really only represented less than half a percent of firms' revenue, which I think is quite telling. If true, that suggests that relative to other kinds of risks that a firm faces --

operational, regulatory, environmental, liability, employment, everything -- cyber may not be such a costly thing for them.

We had done other research asking consumers, how do you feel about firms' behavior in response to these data breaches? Are you happy? Are you not happy? And for the most part, they didn't -- they didn't seem to object. They were relatively happy with firms' responses, what these letters looked like, the kinds of information that were included, and the suggestions. No, it's not great, right? These disclosure laws can only do so much, and the notices can only do so much, but the point is that consumers were not objecting as much as we thought they were. The customer attrition was not as much as we thought they were. So the point is that if the costs to firms aren't as great as we think they are, and if consumers aren't really as mad as we think they are, then what is the incentive for firms to adopt or to improve their practices?

And I would argue maybe that they're doing just the efficient amount. Maybe they are investing as much as they should in order to minimize their costs. Maybe not as much as consumers would want them and security advocates and other privacy advocates,

but maybe they are doing the appropriate amount.

And that had led me to other work on cyber insurance, which I can talk about maybe, but I see that time is up, so I'll stop.

MR. SMITH: Thanks, Sasha.

All right. Well, Rahul Telang -- I'm sorry. So your last name -- oh, good. Okay, good. Rahul Telang is a professor of information systems at Carnegie Mellon University. His research interests lie in two major domains. One is the digital media industry with a particular focus on the economic consequences of the digitization of songs, movies, TV, and books. His second area of work is on the economics of information security and privacy. He's examined the issue of vendors' incentives to improve the quality of their products and the role of policymaking and standards and changing these incentives.

His earlier work explores the challenges of vulnerability disclosure and how competition and policymaking affect these patch release decisions. Recently, he is examining the role of data breach disclosure laws and identity thefts. He was the recipient of an NSF career award for his work on the economics of information security.

So...

MR. TELANG: Thank you. Thank you for having us. So, you know, broadly in the economics of information security and privacy, I'm very interested in trying to understand the firms' incentives and then particularly trying to understand how the market structure -- you know, the market frictions that information transparency actually affect both the firms' incentives to do the right thing -- and we'll define what the right thing is -- and even the consumers' incentives.

So some of my earlier work tried to look at why do software vendors create buggy products and what are the welfare implications of that. And currently that I'm interested in just looking at the data breaches broadly. And I'm just using data breaches as a proxy because actually getting data on the firms' security posture and how much they're investing and where they're investing is just very difficult, not that we should not go after that. It's just that that sort of information is much harder to get.

So we looked at, you know, the hospital industry and tried to understand do hospitals in the competitive markets actually do a better job of investing in security or having fewer data breaches.

You know, the IO market for hospitals is very well-defined. People study about the hospital competition and outcomes very well, so, you know, you can kind of borrow from that literature heavily.

But one other thing that we find is it's not clear at all that the competitive markets are less likely to see fewer -- are more or less likely to see breaches. In fact, we find that it really makes very little difference. And one other thing that we find is that in a setting like hospital, data breaches and information security is the last thing users care about when they're choosing hospitals. If anything, the hospital -- the users care about how nice the building is and what the surgeon is and whether they have all this equipment. And that just means that information security is not one of the features that a hospital can sell in the market and be able to get demand or be able to try to get higher prices. That has interesting implications about, you know, what is the role of policymakers now because the markets may not necessarily create the sort of incentives.

You know, some of the other projects that I'm looking at is at the consumer level, that do consumers actually respond to data breaches. And, you know, one other goal is to actually get the actual

user data, so we are working with a financial institution. And, you know, there are two things that we are noticing. Number one, if there is a direct harm, that is, if the user knows they actually -- there was a direct harm because of some security incidents, we find evidence that they actually punish the bank. So they take their business somewhere else over six to one-year period, with a three to four -- three to five percentage point increase in consumer churn.

But on the other hand, we are also finding to an extent -- and it's very preliminary -- when there is no direct harm, so a retailer got breached, I transact with the retailer but there is no evidence of a direct loss to me, we are finding very little evidence that consumers are willing to punish the retailer. So the longtime impact of the data breach at the retailer seems, at least in our data, you know, very minimal.

And then there's another piece, which is sort of more privacy side that I'll just mention and then pass it on. We are working with -- we ran some randomized experiment on online advertisement. The goal is that the online -- you know, the online ad platforms are using extensive amount of behavioral

data to target us, you know, programmatic advertising, algorithmic advertising. So they're essentially doing selection. They're trying to find people who are more likely to buy and then serve the ad, which is fine as long as the people who are more likely to buy are also responding more to the ads, then that's probably at least somewhat of a win/win.

You know, we might still care about our privacy, bvhrsput at least the advertiser and the ad platforms are better off. And, you know, basically what is our research showing is that that's not true at all, at least in a series of experiment, people who are more likely to buy, and we can see they are more likely to buy from the behavioral data that we have access to, are not necessarily the people who are responding more to ad either. So what we are finding is that the ad platform have all the incentives to target and select people, and they go back and report to advertisers look how good my ad campaign is. It's not clear that the advertisers are necessarily benefitting from paying premium for this very extensive targeting.

So look forward to the discussion.

MR. SMITH: All right, great. Thanks.

So our last panelist is Liad Wagman. Liad

Wagman is an associate professor of economics at Illinois Institute of Technology's Stuart School of Business and visiting associate professor of executive education at Northwestern University's Kellogg School of Management.

He works in the areas of information economics, industrial organization, and entrepreneurship. His focus is on issues of privacy, information utilization and trade, and innovation. His recent works include a study of privacy in financial markets, a study on the tradeoffs associated with increased security via government surveillance, and studies of privacy in oligopolistic markets. And those studies incorporate data access and information consolidation as factors in antitrust considerations.

MR. WAGMAN: All right, so maybe I'll talk a little bit about the privacy aspect that us economists are more used to, that is in the context of price discrimination. And I started my work on this in a context-agnostic way by just looking at the standard models we use like Cournot or Bertrand and so forth. And I found that the impact of whether there is privacy or there isn't privacy on a consumer surplus is not obvious. It's not monotonic. Some privacy is good; too much is bad.

153

I then looked at individual consumers and individual firms in a market, and I saw that the effect on them is also not obvious. It depends on the model we use. It depends on the market structure in question and on the specific context. That means that even in a given market over time some may benefit and some may lose from privacy. And, so, privacy regulation should not be a static thing. It needs to be adjusted dynamically.

I then looked at more context-specific cases. I looked at privacy and financial markets, specifically mortgages. And I looked at the information we disclose as part of our mortgage application process and whether that information can be sold or not. And I found that when it cannot be sold, when we have some degree of privacy there, that prices tend to go up, i.e., mortgage rates. And when they go up, firms have less incentive to screen away consumers. And so standards decrease, and foreclosures might increase, and denial rates decrease. So that's one context-specific study in financial markets.

I also looked at cases of antitrust and whether privacy or lack of can tip the scales one way or another. And what I found is that when firms have

154

consumer data, wide access to it, where they can price discriminate very well, that it's easier to prove antitrust cases that are marginal, that are just on the bounds between being rejected and proved, meaning lack of privacy can intensify competition, which can be good for consumers. So, again, the relationship between privacy and consumer surplus is not obvious. This is not taken into account, any intrinsic value of privacy or issues of data security.

I then looked at cases of government surveillance, something that might not be a popular topic here, but I think it's important to note that even there the relationship between the number of persons intercepted through wire tapping specific to the narcotics-related cases, which is the vast majority of them, and the number of persons that are arrested or convicted is not linear, it's not monotonic.

And I looked at where states are and where the Federal Government in terms of law enforcement is on this nonmonotonic curve. And I found that it's actually -- if you consider it as a Laffer curve, kind of a U shape, it's on the left side of the curve, which is good news.

And another interesting context-specific

155

privacy issue I looked at is physical privacy in a neighborhood. I looked at the effect of short-term rentals in a neighborhood which you might argue hurts neighborhood cohesion and maybe hurts physical privacy around your home. I looked at the effect that short-term rentals have on real estate prices, and by proxy on your physical privacy. And I found that some of it doesn't hurt it but too much does. So, again, there is a nonlinear relationship between the effect and whether you have privacy or not.

So -- and to sum, privacy is hard. It's not easy. There's a lot of aspects to it. If you just look at data privacy, there is data that is collected, there's data that is used, there's data is stored, and there's data that is transmitted. And each of these steps involves privacy considerations.

MR. SMITH: All right, thanks very much.

Great. I guess we'll just get to general questions. So I think, you know, several of you guys sort of touched on the question of what are firms' incentives and how will it balance sort of with efficient outcomes. So, Frank, I think you kind of said, well, they're not doing a lot of things they could be fairly easily doing. So I have a question about that, which is when you say this is low-hanging

156

fruit, like how low-hanging is this?

MR. NAGLE: That's a good question because much like Liad was just talking, it depends on the firm and it depends on the industry, right? So for small firms, low-hanging fruit is actually high-hanging fruit, right? So, you know, you think about mom-and-pop, you know, pizza chain -- pizza restaurant or something like that. For them, investing in somebody to come in and do a security analysis and put in a firewall and all these types of things could be very expensive. For large companies, things like good password policies, closing ports, patching vulnerabilities, you know, those types of things, they're still an investment, but they're comparatively much cheaper.

And, so, you know, something like Equifax, the breach that's in the news now, that was a known vulnerability that was, everybody knew it was a bad thing and should be patched, and it had been gone unpatched at Equifax for at least two or three months. And, so, that -- you know, is that free to fix? No. But is it much cheaper than investing in, you know, a thousand cyber agents to kind of come and help you out and protect your whole company? That's a pretty straightforward thing to invest in.

157

MR. SMITH: Okay, thanks.
And, then, Sasha, you sort of suggested that consumers don't care that much? Is that sort of a fair way to put it? Or --
MR. ROMANOSKY: Yeah, I mean, this has -- this has been the story for a while. It's kind of -- you know, I mean, it's an old problem, right? If -- you know, if consumers really did care, then the firms would start competing on privacy. They would start competing on security. And have we really seen that? I haven't seen much evidence of that.
There may be some instances in sort of niche examples with browsers and certainly products like Tor for anonymizing web traffic, web activity, have increased in popularity, but, you know, I don't think there's anything across the board that would suggest that.
What was the other question?
MR. SMITH: Well, that was basically the question. But I guess one thing I wanted to ask you about in terms of that as well is, I mean, do you have any sense, sort of is this because they really -- they don't think the outcomes are a big deal, or is it because they just sort of don't know how to effectuate a different outcome?

158

MR. ROMANOSKY: Yeah, I mean, again, a well-studied area, and you could -- I mean, I think it's a lot of reasons, right? We enjoy -- we want the benefits now. We can't anticipate the costs later on, right? The costs are intangible. It's very contextual. What I might feel is a privacy invasion, Rahul probably does not and vice versa, right? Change your preferences, change over time, and so the challenge for policymakers, how do you create a, you know, something reasonable, any kind of intervention that can address and can accommodate all these people. I may not like the advertising. You know, somebody else may. And what do you do with that?
But, yeah, I mean, I'll say in the research that we did, asking consumers about their privacy interests and their taste, in their responses, they were -- I wouldn't say quite -- it's not that they were indifferent, and they were, in fact, generally quite positive to firm -- to firm practices.
And, again, if that really is the case and you found some examples of consumer attrition and churn, how they report, how their industry reports, have found a little bit here and little bit there, but, look, if there's nothing driving it, right, if we as a community, if we as consumers are not driving

159

that, then, you know, why should a firm respond to that?
MR. SMITH: Yeah, and so you alluded to Rahul found this evidence that people's reactions are different under different circumstances, right, so can you unpack a little bit more on this idea that the people react more strongly when they sort of see the direct effects?
MR. TELANG: So I think the problem with security or even privacy is that in many, many industry, it's like one of the feature. You know, you go to Home Depot or you go to Walmart, you know, maybe the prices really dominate, you know, your decision-making process, whether Walmart has a good data control policies are something probably I can -- I'm sure most of the people don't care until there is a big breach and then maybe we pay a little bit of attention at that something.
But for some other industry, I think like for financial, like my bank, maybe we really care about it, that is -- are they protecting my information. And we probably pay a little bit more attention. So maybe the data seems to point that, like when there is a breach at the bank or if I'm losing -- I can see in my account that there is a $200

160

harm, I worry about it, even though actually in our data we find the bank actually compensates me, so I get my money back. Even then, I feel -- at least the data seems to suggest that people are a little concerned that why is this fraud perpetrated on my account. Why didn't bank do enough or wasn't it proactive enough?
So for some industry, we feel like security is an important feature. I think at least the users feel it's an important feature. I think financial institutions are probably a good example. But I think for a retailer or hospitals or some other thing, I'm not so sure that consumers -- and can't blame them either -- you know, maybe that'd be default. We expect them to have it, and that's not something that's going to drive the marketing or the pricing decision, or you can charge premium for that, let's put it this way, but maybe the -- so I think that's really what we seem to find, and it's probably consistent with what users should rationally behave.
MR. SMITH: So, actually, you know, you raise -- so one of the questions I guess I have about that is to what extent when you're asking -- when you're looking at this question, to what extent do consumers seem to understand sort of what differences

161

there might be or are they pretty indifferent? Are they knowledgeable or --

MR. TELANG: Yeah, so, in -- I think that Sasha's probably had this survey that he did where he asked people about -- or at least the research tried to ask people about their perception. You know, it's -- my research, we actually had the actual behavior, but we didn't actually ask them about their perception. My suspicion is that it's kind of correlated, which is for the retailer, if there's a breach, they pay a little bit of attention and then kind of ignore, you know, the future transaction when they make the decision. For probably financial institution and bank where we keep our sensitive information, I think people not only behaviorally show that they care about it, but perceptually they probably care about it. That would be my, I think, sensible conjecture.

MR. SMITH: Liad, I'm going to shift the topic a little bit, but sort of still getting to the question sort of -- sort of efficiency. You know, you talked a lot about the sort of idea of there can be too little privacy and too much privacy. Is there any sort of way to think about when we might expect that to be, you know, on either side?

162

MR. WAGMAN: Right. So let me give an example. If we think about, for example, being able to maintain your privacy at some cost, and if we imagine this cost as something continuous that a regulator can control, the finding is that when this cost is too low, consumers end up being harmed and firms end up being harmed. And when this cost is too high, firms are actually happier, consumers not so much.

Now, once you engage in repeated interaction between firms and consumers, these findings change. Firms, in fact, might want to commit to a level of privacy because they'll be able to retain consumers over repeated interactions. So what we find is that even in these repeated interactions, having too much privacy or too little privacy ends up being bad for consumers. And the reason is too little privacy, consumers, at least the lower willingness to pay consumers, don't get the benefits of price discrimination.

And when privacy is too expensive or too hard, then firms don't need to try to give reasons for consumers to be tracked to give their information. So somewhere in the middle kind of gives a sweet spot where consumers are willing to give the information

163

and willing to be tracked in order to get the benefits, and firms are happy as well because they're able to price discriminate.

MR. SMITH: So when you say the middle, is that sort of an exogenous dimension of cost or is that something --

MR. WAGMAN: No, so, again, it's market-specific, it's industry-specific.

MR. SMITH: Mm-hmm.

MR. WAGMAN: And, you know, this is what makes it hard. Now, even in a particular market, among consumers, there are going to be winners and losers. There are going to be some who are happy that there is privacy or that there isn't privacy. And those groups of individuals might change, depending on market structure, which itself can change over time. So it's a dynamic question of what's efficient.

MR. ROMANOSKY: Can I add one thing?

MR. SMITH: Absolutely.

MR. ROMANOSKY: So I think that leads to a really interesting question, which is whether or not privacy regulations, say state laws, actually harm consumers or not, are actually in their best interest or not.

And, so, one way you might think of that is

164

if we want to define privacy as the control over our information, the right to control, the ability to control our information, say in a financial setting, you might wonder about what the effect -- so let's say there were state-level laws that allowed for more or less sharing of financial information amongst financial institutions. So some states were very strict and required and permitted very little sharing of information between financial institutions; other states were very permissive in the sharing.

And, so, the question is more or less is information-sharing better or worse for consumers. And, so, privacy advocates would certainly argue that, no, I want control over my information, I don't want that to be shared amongst financial institutions. On the other hand, what that might lead to is higher price of credit, right? So the less information the bank has about you, the less they're able to assess your financial risk, the more likely they're going to charge you -- the more they're going to charge you higher rates for borrowing money.

MR. SMITH: Right.

MR. ROMANOSKY: And, so, I don't know if that -- I'm not saying that that's true. I'm just saying that that's a reasonable question, and that's a

165

testable question. And I think there are all kinds of examples of state-level -- you know, state-level laws make for a great sort of empirical study, but I think there are all kinds of different ways that we could start to think about how state-level variation in different kinds of privacy laws drive different kinds of outcomes, and maybe those that we wouldn't -- we would actually consider.

There hasn't been a lot of work in that area, but I think if there's really any opportunity to pursue that it's -- it would be a dynamite thing to be able to show.

MR. WAGMAN: So we were able to test something along these lines.

MR. ROMANOSKY: Okay, there's one great paper on that.

MR. WAGMAN: In the Bay Area, where some counties adopted stricter privacy financial laws and some did not. So we had a control group, and we were able to test this. And we found that there was an effect. There was a significant effect where once privacy was enforced in those counties, their rates on mortgages increased. They were charging higher prices, and they were approving more mortgages. So their denial rates decreased. Their standards

166

decreased.

MR. TELANG: I mean, don't you feel that if people are heterogeneous in their preference for privacy or whatever you would think that the firms would also be heterogeneous in terms of providing that? So some firms would say, okay, you want a lot of privacy, here it is, and if you don't want too much, then here it is? But we don't see a whole lot of that happening. It could be because of the earlier talk about the concentration. Maybe there are some -- you know, we can't live without Facebook, and there is really no competition to it possible because of all these effects, so then whatever is the Facebook privacy policy is really what we have to live with.

But you would think, right, I mean, if I have heterogeneous preferences, and if I can market it to consumers, that, you know, this is what my privacy is or this is my security, you would think that the market should sort itself somewhat without FTC or policymakers intervening too much.

MR. WAGMAN: Right. I think this is where, you know, privacy is a second-order effect. It comes in, and consumers usually treat price as the, you know, the driving factor. And privacy just comes as a second-order effect, a second-order consideration.

167

And, so, even if it's different, you know, it might not be what drives consumers to the product. And we see this in the market for mortgages with different privacy policies and different mortgage rates.

MR. NAGLE: And there are some implications and kind of externalities beyond just the rates and things like that. There's a study by Catherine Tucker and some friends that looks at the impact on innovation and the ability of the firms to innovate and actually shows that increased privacy slows down innovation. And, so, if we think of innovation as probably a good thing, then this balance and kind of the sweet spot of regulation also factors in beyond just the individuals but to the ability of firms to innovate as well.

MR. SMITH: Great. I don't know if there are any sort of more general thoughts on the question of sort of efficiency and privacy as it happens in the market.

MR. WAGMAN: So just a quick thought. I mean, Sasha mentioned that a lot of privacy considerations among consumers involves some form of regret where you give your information away and then you realize later, oh, what did I do. But most regulators' guidelines, at least, pertain to ex ante

168

consent. Give my consent now or not. Almost none talk about ex post consent, where my information is already out there and I want it withdrawn. There have been some, you know, policy experiments in the EU along these lines but not much in the U.S. So, you know, that would be interesting to explore.

MR. TELANG: So if I take it in a slightly different direction, I think we are very interested in understanding are firms doing -- you know, investing optimally in security. You know, you don't want them to spend too much. You don't want them to spend too little. And, you know, what is the ROI and everything. But sometimes I feel that this can get very complicated if your adversary is some state actor.

So suppose you are being attacked by somebody in some other country who might have very nonmonetary incentives to actually -- so they want to attack you because -- not because they want to steal your data and make money off it. They just want to have a -- cause a significant reputational damage to you.

In this situation, it's a little -- it's very challenging to think about the private investment by a firm would be the right strategy to fight

169

against, something that's happening. Then you kind of go into this, you know, is there a role for government here, is there a role for some public investment, whether it's diplomatically or whether any -- and it just opens a can of worms.

But it also means that the whole, you know, modeling gets very complicated because, you know, what are you modeling? You know, are you -- what exactly is your model of investing in security when, you know, you have some actors which are probably not driven by economics alone.

MR. NAGLE: And along those lines as well, these state-sponsored actors often will attack even small companies that have -- you know, they're not going after them at all, but they want some IP address in the U.S. to base their next attack against the bigger company or the better target or whoever. And, so, even if we think about the small places that have, you know, limited kind of juicy data or juicy whatever that they want to steal, they're still kind of getting caught up in these super-high, you know, priced kind of attacks, right? The super-expensive attack.

MR. ROMANOSKY: Yeah, and I'd -- I mean, I'd reiterate that it's still an outstanding question, right? It's one that's plagued the industry for

170

decades of how much should firms invest and are they investing optimally and how would you even know. And, again, privacy and security advocates would argue that, no, firms are not investing because look at all these breaches that occur. And I would argue that that's not evident, that they're not investing optimally, at least for their own interest. Even if you take Target and Equifax and even if the cost is $100 million, that's still not evidence that they're not investing optimally.

The other question that we still don't know is what kinds of security controls matter and by how much, right? We could all think of different kinds of technologies to implement that we would think would reduce risk of any given firm by a certain amount, but, I mean, even I can't tell you with all the experience that I have of by how much that should reduce a firm's or increase a firm's security posture. We just don't know.

The one place that I think we could answer that is with insurance. So any given firm, right, you would need to know this marginal benefit, the marginal cost in order to assess this. They don't really operate that way. Even a government agency doesn't really have that information. But insurance

171

companies, right, through their claims data, they collected this body of data and observe these incidents. They have information about the firms ostensibly. They have information about the different kinds of security controls because of the applications that they provide to the applicant in order to sign up for the policy.

And, so, there is a potential there for -- you know, I mean, it's not very complicated, it just runs through a regression to understand what the marginal effect is of different kinds of controls in preventing a claim and a breach and even therefore to understand the relative effects of one versus the other. And that's a dynamite thing to be able to do. I haven't encountered any firm, any carrier, that's doing anything like that, but it's possible to do it, and I look forward to the day when they start to do that.

MR. TELANG: It's the IT productivity question, for a long time we had no clue, then we started collecting good data. Maybe something good was going to happen. I'm not very optimistic because this was a question when I was doing doctorate dissertation. I'm glad I didn't attack it. But it's one of those things where we don't have good measures

172

at all and, you know, both the technology, organization, management, they interact in ways that's really, really hard to predict good econometrics. You see case studies and anecdotes, but that's not -- I don't think that's very convincing.

MR. NAGLE: And, actually, to echo that, a few years after your dissertation I thought that was going to be my dissertation question as well, and it turned out there was no data and I couldn't do anything about it. Although I want to get back to the insurance angle is interesting because as we all know, insurance changes incentives and behavior as well, right? So are these companies -- if they know they get hosed by a bad guy that they -- the insurance company is going to clean up and take care of the loss? Does that mean -- lead them to underinvest in security? I'm not sure, but there might be data to chip away at that.

MR. ROMANOSKY: Yeah, I mean, there are all kinds -- I mean, this is why I've been studying it for a while. I've been trying to get at these questions of, you know, does the insurance even improve incentives, right? That's an outstanding question. We don't know. In theory, it's a testable one. I don't have data on the adoption time of any given

173

company, but if we had it, it would be -- it would be answerable.

Yeah, does it lead -- I mean, the same, you know, information asymmetries that exist with any kind of insurance, you know, may still exist. If I'm a firm, I can buy insurance or I can invest. Why do I need to do both, right? Does that occur and to what extent?

Is it true or, you know, how much information do the carriers need in order to create the right incentives for firms to improve? Right, I think that gets back to they need to understand what kinds of security controls matter. So should they incentivize firewalls, two-factor authentication, better encryption, cloud services, et cetera? From what I've seen, they don't know that. They don't have the answers, right?

I've seen the price schedules. I've seen exactly the variables that they use to price the premiums and the effects on the premiums, like the -- I mean, it's a linear product of a bunch of different variables, right, so I can see if some carriers feel that if accounting firms pose a lower risk so they have a multiplier of .85 versus government agencies are a higher risk and you multiply by 1.2, for

174

example. But they still don't really have a good feeling for how to craft those and good justifications for any of those numbers. But I think that will just improve over time.

MR. SMITH: So, yeah, so that gets kind of into a product question about sort of what are some of the things that we still really need to figure out in this area. What are the big questions?

MR. WAGMAN: So from the perspective of privacy, I think there's been very little link in the literature between privacy and security, right? They've mostly been studied separately, and I'm partially guilty of the same thing. Having worked on a survey of the literature recently, I tried to tie them together, and I think there's a lot more that can be done there. So I think there's great opportunity for theoretical and empirical research to try to tie them together.

MR. SMITH: Can you talk a little bit like what that would look like, or --

MR. WAGMAN: Right. So I think I indicated earlier that privacy, at least economists have looked at it in IO is mostly revolved around price discrimination or search and seizure. And that's quite limited because there's this privacy in a bunch

175

of other things as well. There's privacy in data storage and data transmission. Data that is stored in itself can be made more private by anonymizing it and so forth. And I think these considerations have largely been ignored, at least in the economics angle. Some computer scientists have looked at it, but not many economists. And I think there's a lot of opportunity there.

MR. TELANG: I think -- and, you know, many people have thought and commented on it, but when it comes to security particularly and privacy for sure as well is like, can we even say that there's a market failure? And what are the dimensions of those market failures? What are the things that are leading to these market failures?

Then we can ask the question, what is the good policy intervention. And then how effective those policy interventions are, right? I mean, the data breach notification law was passed, what, 10, 15 years ago now? It's been around, and I don't think that even now we understand, you know, if you talk to the industry people, they'll come and say it's a lot of -- a bunch of checkmarks that I have to do, and I don't know what I get in return. Or they say it's so sometimes outdated that we actually do a whole lot

176

better than what some of these laws are telling us to do.

So sometimes you hear from firms that it's very costly and onerous, but then you look at what the benefits are and then you're back to sort of square one. Some of it is just because the observation nature of data makes it so difficult to do any sort of, you know, sensible identification. You can't run a randomized experiment here. There's really no good exogenous shifter.

But those are the fundamentals, I think, you know, we don't know at some level where the market is failing. Or even if we know, we don't know what sort of policies would make sense and then come back in a while away, you know, is this the right policy? Can we tweak it? What way we should be tweaking it? So I think there are a lot of interesting questions both at the macro as well as at the micro level.

MR. NAGLE: And to kind of add on to something that's been underlying all of this is that, again, a lot of these things are difficult to price, right? What's the value of your Social Security number? And your Social Security number being safe, right? We don't know. And in the case of a lot of the firms that we used to do investigations of, a lot of it was related to intellectual property, right?

So a large multinational conglomerate got broken into; intellectual property was stolen for a widget; and that widget shows up on the black market for, you know, a third of the cost that it actually costs, you know, the company to make. And they end up shutting down this entire business unit, right? So they got breached, and they shut down the business unit and all the future profits that might stem from that.

And, so, how do you kind of value that as well in terms of it's just intellectual property, right? It's an idea. We know it has value, but how do you actually put a future number on that so you know how much to invest in protecting that idea?

MR. SMITH: Okay, so, I think it sounds like there's a lot of sort of sense that we don't really know what the market failures are or where there should be policy interventions. Are there any thoughts about sort of what government might do in the short term in terms of thinking about policy -- towards privacy data and security?

MR. NAGLE: One thing I always think of just is a pure awareness, right? So educating the population, and this is one thing that is known to

work fairly well in firm context. Presumably it would also work reasonably well in the broader populace, but everybody wants to invest their security dollars in the newest, latest, greatest technology to actually, you know, prevent the breach, right?

But how do most -- or a lot of breaches happen now is somebody clicks on an email that has a bad link and then bad things happen, right? So educating, you know, the employees, but also the general populace that this stuff is going on may be a cost-effective way to at least start approaching this.

MR. WAGMAN: I would add to that that the Government did step in in financial markets, for example, and made privacy disclosures very easy to read. It's basically a table that you can quickly go through. And, so, you know, it improves awareness, it improves understanding. I think there's very little of that in other markets, and that would go a long way.

MR. TELANG: So I feel like, sure, we cannot stop the data breaches, but I think we can do a whole lot more to control the cost that happens post data breach. So I think Equifax being a good example, right? Probably the breach itself was bad, but the response itself was so sort of incompetent that you

would be, like, why would you do that. I mean, you know, you're -- we are in 2017, we should be expecting to be breached. And conditional on breach, we ought to have some sensible plan so that we make sure that the damage is contained.

So maybe there is some -- maybe there is some role for policymakers to say, okay, you know, sure, you know, you got breached, we give you benefit of doubt, but you really have no benefit of doubt on how you respond to the breach. I mean, there has to be some way. So containing the damage is something I think we should probably be focusing on, rather than saying how much dollars to spend and reduce the breach and that it should be zero probably. Probably that will never happen, but I think we can do a lot more in making sure.

In fact, how much consumer is harmed itself is not clear. Okay, there's a breach, hundred million records got breached, but so what? I mean, like, what does that mean, right? I mean...

MR. SMITH: So is there a sort of a practical set of things that firms should do when there's a breach? Is that, like, a pretty clear answer?

MR. NAGLE: There's, like, the industry

standards of you have your team that -- your response team that includes not only the techies but also legal, also marketing and PR because you're going to have to, you know, publicize what you're doing and kind of, you know, you're supposed to have your strike team on speed dial, right? And, so, there are kind of standard sets of best practices pre-breach that help you know what to do so you're not running around in a panic. And I agree, Equifax's response was certainly not as good as it should have been.

MR. SMITH: So does this dovetail a little with Liad's point about, you know, how much data do you really need kind of issues? Is that sort of a similar feel in terms of we know things are going to happen, so let's minimize?

MR. WAGMAN: I think with the way the incentives are set up now firms want to collect as much as possible because the data itself often is the product or part of the product. And you don't know what you're going to need tomorrow. So the way the incentives are set up now, firms want to store more and more. So, you know, it's --

MR. NAGLE: Which, of course, makes it much worse when a breach inevitably happens, right?

MR. SMITH: But is that a market failure, or is it just that's --

MR. WAGMAN: I don't want to step on Rahul's toes here, but, you know, I think tastes for privacy have -- are constantly shifting, right? Things that used to be punishment, for example, you'd be put on some registry and public records, now people voluntarily want to be on some sort of public record, right, whether it's Facebook or other social media. So tastes are fluctuating, so it's hard to pinpoint the failure, but if firms are overstoring data, it can be showed in simple theoretical models that this is an inefficiency.

MR. TELANG: I don't know what government can say and tell firms what to store and what not to, so that is a -- that's really being -- you know, I don't think it will work at all. You can only think about the consequences that if you were to lose what are the consequences. And those carrots and sticks have to be in place to encourage them to do the right thing around what data they should have and what data they shouldn't have. I think that would be probably a more practical and implementable strategy versus kind of dictating or even saying anything that how much data you are to store.

MR. SMITH: So more time is outcomes in some

sense.

MR. TELANG: I think so.

MR. SMITH: So we have a clock that's counting down. I don't totally know what it corresponds to. I think it corresponds to in 20 seconds it's time to ask the audience questions -- or open up the -- for audience questions. So why don't we just move to that. So, yeah, any questions from the audience?

Nathan.

MR. WILSON: So there were multiple references to us not knowing the costs of breaches. Can't we at least establish some sort of lower bound by looking at how a breach correlates with the incidence of, you know, stolen identities and then there's -- I presume there must be some estimate of the -- you know, the hours spent dealing with that plus potentially some expenditures. Or is that data --

MR. ROMANOSKY: I don't know if we say we don't know the cost of breaches. So I actually have a paper on the cost of breaches, and it turns out to not be as high as we think it is. So the typical industry reports are in millions of dollars -- $4, $5, $6 million, and the reason they're high is because they

report -- what they're trying to report is the typical, right? So they report the mean. But because loss distribution is so skewed, that's a pretty poor representation. We look at the median, which is much less, a couple hundred thousand dollars.

MR. SMITH: Sasha, are you talking about the cost to the firm, the cost to consumers, or --

MR. ROMANOSKY: Right, sorry, the cost to the firms, strictly to the firm.

MR. SMITH: Yeah, I think Nathan's question was maybe more about --

MR. WILSON: Right.

MR. SMITH: -- consumers, cost to consumers.

MR. WILSON: Does the cost to the firm -- how does that compare to the inferred cost to the populace or the affected populace?

MR. ROMANOSKY: Oh, yeah, I'm sorry. Right, and so the reports -- right, the reports are scattered. Bureau of Justice Statistics has had some -- it's kind of sporadic over a few years. They've tried to collect those data, and, again, it's still very skewed. And the median might be close to zero, right? But for those people that did report losses, it was in the hundreds of dollars, right?

Now, it's -- right, this is always the

problem with privacy, right? And between all of us here, it's one of the reasons why I avoid privacy research, it's just because it's so squishy and nebulous and difficult to figure out, right? So, you know, one measure of the harm, the privacy harm is looking at the dollars lost, but if it's true that the banks -- and a lot of these are due to financial fraud -- if the banks are always covering your costs, then really the harm is zero. But that's not really the extent of it because there are lots of emotional distress, and certainly, you know, very legitimate kinds of severe kinds of, like, forms of identity theft. And, so, I'm not to discount those, but relatively minor in terms of numbers.

And, so, how do you put all of that together? How do you mash it all together in some kind of metric that is sort of useful for us as researchers or for policymakers or for anyone to try and figure out. I don't have an answer for that.

MR. NAGLE: And along -- to go a little further, it also depends on what is stolen, right? So credit cards, absolutely, the bank makes you whole, not a big deal. Intellectual property, if you're a company, harder. Once it's out there, it's out there. So there are -- you know, you shut down a business line or have different responses.

And then somewhere in the middle is kind of the Equifax breach, right? So very easy to change a credit card number, very hard to change your birthday, right? That's pretty much there from when you're born. And, so, you can't change that once that's out in the open. And, so, what is the cost there? It's a little bit different depending on what data is stolen.

MR. TELANG: If you're willing to -- if you're willing to make an assumption that suppose after the breach, if you give me credit freeze and credit monitoring service, then I will not be harmed, then you can kind of look at that as saying, okay, you know, this is worth $100, I have to service, you know $100 million or 100 million consumers, maybe you can kind of get ballpark numbers, but as I said, how do I value the identity theft that happens two years after the breach happened?

MR. WAGMAN: So maybe you can do some different studies on increases in identity theft after breaches, and I would assume that a lot of people don't take advantage of credit monitoring or credit freezes when they're offered, especially when the source that offers them doesn't seem very reliable.

MR. NAGLE: And the source -- to get the

free credit monitoring, you have to sign away your right to sue the source.

MR. WAGMAN: Right, exactly.

MR. NAGLE: Which skews my incentives.

MR. WILSON: Thanks.

MR. SMITH: I'll just make a quick note that on December 12th we're having a workshop on informational injuries, which part of the goal is to try to get some new thoughts on how to sort of think about measuring these kind of harms and even just conceptualizing them. So if anyone's interested, that will be happening, I guess, next month.

MS. JIN: Yeah, I really enjoyed the panel. So here is a question. I know variation in state regulation is great for research, but a lot of company names we heard today like Target, Home Depot, or Google or Facebook, they all operate in many, many states. So how relevant is local regulation in this?

MR. WAGMAN: In the case of financial markets, there's a national benchmark based on the Gramm-Leach-Bliley Act, and then local areas can put stricter regulations in place.

MR. NAGLE: And there's also -- and I think it's California, a lot of their regulations are written so that if you do business in California,

which all these companies, you know, Facebook, Google, everybody does something, even if they're not located there, then for all of the customers, they have to kind of hit the bar, right? So for -- if California is moving the bar up, then everybody else benefits, at least in the U.S.

MR. TELANG: This is what I don't like about policymakers. They create policy but just make it so hard to do any identification. You're exactly right. I mean, you know, if they affect how independent your observations are, it's --

MR. ROMANOSKY: You might get struck by lightning as you leave the building. But there's lots of -- there are lots of different kinds of privacy laws, right? Local, DMV-related privacy laws, you know, nursing privacy laws, surveillance privacy laws, blood type privacy laws, which are all very localized to the state level. There's lots of variation there, and a couple of people have done -- written some compendiums of these state laws and put them together and tracked them over the years. And it's great stuff.

The trouble is finding the outcomes that are used mentioning -- that are useful to measure and to try to associate the two and come up with kind of a

useful paper on that to try and answer a good question, but there are certainly lots of different kinds. And, yeah, the point about the breach laws is well taken, right? If you do business in that state, then, I mean, you know it well. But there's lots to go on, the trick is -- that I found is trying to find a useful outcome measure to study.

MR. WAGMAN: I think for those large firms, it's easier to cope with a patchwork of laws, but it might stifle innovation in the sense that a small company, you know, it would be really hard.

MR. SMITH: Okay, I think -- I want to thank everyone for their patience and thank the panelists for a really, really interesting conversation. And that's it for this panel. Thanks so much, guys.

(Applause.)

MR. ROSENBAUM: I just want to give a final thank-you to all of you for coming and to everyone who presented and helped to facilitate the conference. I hope to see you next year. Have a good weekend.

(Conference adjourned at 1:21 p.m.)

CERTIFICATE OF REPORTER

I, Jennifer Metcalf Razzino, 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.

s/Jennifer M. Razzino
JENNIFER M. RAZZINO, CER

L