In the Matter of:

The Eleventh Annual FTC Microeconomics Conference

November 1, 2018
First Version

Condensed Transcript with Word Index

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

UNITED STATES FEDERAL TRADE COMMISSION

THE ELEVENTH ANNUAL
FEDERAL TRADE COMMISSION MICROECONOMICS CONFERENCE

Federal Trade Commission
FTC Constitution Center
400-7th Street, S.W.
Washington, D.C.

Thursday, November 1, 2018
9:00 a.m.

Sponsored by:
Federal Trade Commission Bureau of Economics
and
Northwestern University's Searle Center on Law,
Regulation and Economic Growth

C O N T E N T S

WELCOMING REMARKS

- - - - - -

MR. VITA: Okay, let's get started, everybody. Good morning. My name is Mike Vita. I'm the Deputy Director for Research here at the FTC's Bureau of Economics. Thanks to you all for coming, and welcome to the Eleventh Annual FTC Annual Microeconomics Conference, where we attempt to combine cutting-edge academic research with discussions of real-world policy problems. As always, we're grateful to Northwestern University and the Searle Center for their continued cosponsorship of this conference.

For those of you who are from other institutions, a few words about us here at the FTC. As you probably know, the FTC is an independent agency that, along with the Department of Justice, enforces the antitrust laws. Our other major mission here at the FTC is enforcement of federal consumer protection law. These enforcement missions are supported by the FTC's Bureau of Economics, which is a group of about 80 Ph.D. economists, which makes it one of the largest groups of applied microeconomists in the Federal Government.

At the FTC, we believe very strongly that these twin enforcement missions reinforce and complement

each other. Competition, we think, is most effective when consumers are making well-informed choices and decisions, and consumer protection works best when consumers have real alternatives.

Today's conference, like its predecessors, helps ensure that the FTC's actions are informed and guided by the best possible economic analysis. So I think, as we always do, I think we'll have a fantastic conference this year. In addition to the usual cutting-edge papers that we always feature, tomorrow we have a panel discussion on the estimation of markups, a topic that's become pretty important in antitrust circles these days.

Before the first panel starts, just a few acknowledgments and then a few official announcements. First, let me take a moment to thank Ted Rosenbaum, Nathan Wilson, and Alex Avramov of the FTC for their hard work in putting together the conference; Julie Carlson, Antara Dutta, and Nathan Petek for their assistance to the scientific committee; to a large group of BE economists, who I won't mention by name, who gave feedback on the various submissions that we received; and to our scientific committee of academics, David Besanko of Northwestern, Katja Seim of Penn, and Ali Hortaçşu of Chicago.

Ali, if you're in the audience, I hope I pronounced your name right. I got multiple opinions about that. They were all different.

I also want to thank our wonderful BE administrative team, who always do incredible work behind the scenes to ensure that the conference comes off seamlessly, Maria Villaflor, Kevin Richardson, Neal Reed, Constance Herasingh, Priscilla Thompson, and Tammy John.

On that note, on the administrative note, right at the moment we do not have WIFI information for you guests, but we will soon. So that will be -- we will update that.

Then I guess I'm supposed to -- Ted says I have to read some important legalese that we're compelled to talk about. First, silence your mobile phones; I'm sure you've all done that. Please be aware that if you leave Constitution Center -- that's this building -- for any reason during the workshop, you will have to go back through security screening again.

This is in bold. Most of you received a lanyard with a plastic FTC Event security badge. We re-use these, so when you leave for the day, please return your badge. You know, money's tight. We can't replace those.

If an emergency requires that you leave the conference center but remain in the building, follow the instruction provided over the building PA system. If an emergency occurs that requires the evacuation, an alarm will sound, and just follow everybody else.

If you notice any suspicious activity, please alert building security. I don't know if that includes suspicious activity that takes place during the panels, but any other kind of suspicious activity, alert building security.

Lastly, please be advised -- oh, importantly, restrooms are located in the hallway just outside the conference room. And last, please be advised that this event may be photographed, webcast, or recorded. By participating in this event, you are agreeing that your image and anything you say or submit may be posted indefinitely at ftc.gov or on one of the Commission's publicly available social media sites. So it will live forever. So choose your words carefully.

Okay. That, I think, concludes that. So let me just turn it over to Julie Carlson, who will introduce today's first panel. Thank you.

(End of session.)

PAPER SESSION:
PUBLIC COMMUNICATION AND
COLLUSION IN THE AIRLINE INDUSTRY
- - - - -

MS. CARLSON: Welcome. It's my pleasure to open our first session, which was organized by David Besanko of Northwestern. So we will have two papers in this session, and each presenter will have 25 minutes to present, and then after each presentation, we will have a ten-minute discussion, and then we will have about ten minutes left over for questions from the audience.

So our first paper is by Gaurab Aryal from the University of Virginia, who will be presenting Public Communication and Collusion in the Airline Industry.

MR. ARYAL: Thank you. Thanks to the organizer for accepting the paper. This is joint work with Federico, who is at the -- he is also at Virginia, and Ben, who was a grad student but now at Cornell. So this paper is about public communication -- I'll explain exactly what that is -- and how that can facilitate collusion; in particular, the industry that we look at is the airline industry.

So just the big picture. So, you know, the idea that -- so there are two kind of competing

institutions or laws that we look at. One is antitrust, which forbids firms from communicating with each other in order to kind of deter collusion, but on the other hand, you have financial regulations, which tries to promote transparent communication, all right?

So the question that we are interested in is what if the second one helps evade the first one, okay? What if these transparency laws facilitate collusion? This is actually pretty well thought out by the OECD, so it says that information exchanges can offer firms point of coordination of focal points. Of course, these are all abstract terms. What is focal point? What is coordination? So we tried to find an empirical evidence of that in the data.

Ultimately, I mean, of course, we don't address this question in the paper, but ultimately we are interested, of course, as empirical IO-ists is what kind of information should firms be allowed to share in public, okay? And we leave that question as it is.

So the main objective of today's talk and the paper is to ask the following question: Do managers in legacy U.S. airlines use their earnings call to communicate with other legacy airlines in reducing the number of seats sold in the domestic U.S. airline market?

Okay, and in particular, when we talk about communication, we focus on this concept of capacity discipline, okay? So any word related to capacity discipline, that will be what we will call communication, of course. Just to give -- why is information important or communication important? A priori, you think that given, you know, the nature of the business, stochastic demand, kind of difficult to monitor each other, you think that collusion among airlines would be difficult. That's the a priori thought.

But, of course, there's these three really super papers by Yu Awaya and Vijay Krishna, one is in AER, the first one, and what they show is that if you have private monitoring, meaning I can only observe my own action, but you can sense some cheap-talk, some information out, then they show that in many cases you can do better than Nash, okay, meaning you can sustain a collusive outcome.

And so what we are going to do is kind of think about this in sort of like a reduced-form way, is giving -- taking this as a benchmark theory model, we are going to try to see if there is any evidence of this in the airline industry.

So in terms of the data and the methodology,

what do we do? We basically build a data set of public communication. We read through all the earnings calls. And so every quarter, publicly traded companies hold earnings calls where they communicate about the future strategies about the companies, and all of these are transcribed and recorded. So we basically read all of them and try and determine which of the -- which quarter's earning calls were pertinent, pertinent meaning there was some communication about capacity discipline.

And then once we have figured that out, we try to estimate or actually we estimate the causal effect of communication -- I'll explain exactly what that is -- on the number of seats made available in the domestic market, okay?

So there are some issues. So exactly how do we approach that? The first is we -- the first thing we do is we ask, do carriers change capacity after discussing capacity discipline? So think about quarter one, where everybody, all the legacy carriers in a market, a market defined by airport-to-airport pairs, mention capacity discipline. Then we look at, subsequently, does that lead to a reduction in the number of seats being sold in the following quarter? And the answer is yes.

So we find that, on average, the airlines reduce about 1.45 percent, and that's a -- you know, in terms of how big that number is, in general, the average change in capacity is about 3.5. So it's a pretty big number.

The other thing that we have to kind of determine is that is this a collusion or is this just, you know, the airlines using these earning calls to be more transparent about their strategies? And if they are, indeed, just being transparent about the strategies, then the fact that we find a reduction in capacity does not necessarily mean that they are coordinating, okay?

And so if that was the case, then we look at things like, you know, imagine the airlines were the only one who mentioned the word "capacity discipline," while everybody else serving the market do not, then do we see a reduction?

If it was serving the purpose of just being transparent, then we would imagine there to be a reduction, but we find none, okay? And we also then look at other ways; for example, we also look at what if everybody, other than -- except -- so if there are five airlines serving a market, suppose four out of five talk and mention capacity discipline? Do we see

a reduction?

If, indeed, they were just serving their transparency problem, then we should have seen a reduction, but we don't see any reduction, all right? So these are some sort of not direct but kind of indirect evidence that it seems like this capacity or these communications are not serving the purpose that they were supposed to, okay?

The third is we also have to deal with some -- because we are using words and text, there might be some other words that are connected and we miss out on that, so we look at that as well. And the other issue is, as I'll explain in a bit, the way we define communication requires some market structure. So we know from previous literature that market structure can be endogenous, depends on some other unobservable that is not accounted for in the data. So the question is, should we then be concerned about these? And we addressed this.

First, as far as the communication part is concerned, we look at -- we do some conditional exogeneity test, and we find that our result is consistent. And we also do IV, actually control function approach, to try and address the fact that the market structure might be endogenous, and we still

find that there is a significant reduction in capacity whenever they talk about communication, okay? I'll explain all of these in detail.

For the remainder of the talk, I'll just briefly talk about the data. I will not go into the detail. There are two basic parts of the data. One is the transcript data, which is what we've collected from the earnings call, and the second is the usual airline data, which is the T-100 and DB1P. We augment the airline data by buying -- so we augment that with an OAG data, which allows us to take care of the differences between the ticketing carrier and the operating carrier. So it's possible that if two major airlines subcontract to a local one, the T-100 is not going to capture that part, so we know exactly who the regional carrier is contracting with, okay?

I'll talk about the empirical analysis. As I said, I'll address these two or three concerns that I just brought up, and then I'll conclude. This is the data on transcripts. So each row is an airline, so there are 11 airlines. Each column is a quarter. So we start from the last quarter, 2002, all the way to 2016, all these different colors.

The colors that we should focus on is the light green one. So that means that -- say, if I look at

one particular -- so, if -- take Southwest, for example, at the bottom. Quarter four, 2002, it's green. That means that for that quarter we have collected the transcript data, okay? And other colors are different reasons for which we do not have the transcript data.

For example, if they are privately owned, we don't have it. If they're before -- just after merger, we don't have it. And then in some cases, the black ones, we don't know -- we don't have the data, but we don't know why the data is not available. When we do the regression, we try to control for that as well, okay?

So not all of these green ones are the ones where the airlines talk about capacity discipline, okay? So how do we -- just a one-page thing about text to data. So what do we do? Basically we take all these text documents and then we use the natural language processing and we try to flag any quarter in which the word -- by "word," I mean the semantic of "capacity discipline" -- shows up, okay? So it's possible that the airlines do not always use exactly "capacity discipline," but as long as they imply capacity discipline, we pick it up.

We do a bunch of robustness on that. We verify

17

ourselves. We hired an independent RA to read through this completely carefully and give us the data the way the RA thought about it, so we can -- so we double-checked everything. Everything is in the paper.

And in some cases, when it's not absolutely clear what's happening, we read the transcripts ourselves. Three of us read independently, and we try to, you know, see if we all agree that this is pertinent, i.e., this is about capacity discipline, or this is not pertinent, okay?

So to just give you examples of how these words "capacity discipline" crop up, this is U.S. Airways. Main line passenger revenue were up by 2.1 billion, da-da-da, and continued industry capacity discipline. So they are basically trying to say that our revenue went up because there was an industrywide capacity discipline, so everybody were disciplined when they were choosing the capacity.

Or the CEO of Delta, you have heard us consistently state that we must be disciplined with capacity. So as far as we're concerned, even though it's not exactly capacity discipline, the second one, it pertains to the same notion of capacity discipline. So we picked both of these instances.

18

So it's only when -- so imagine a market that is served, let's say, only by U.S. Airways and Delta. In that particular quarter, if we find that both of these -- you know, these airlines are using this capacity discipline, then we say that these guys are communicating. So we need at least two people to communicate, two legacy carriers to communicate, and all of them must be communicating prior, okay? And that's what our definition of "capacity discipline" is, and that's going to be important.

And as I said, now, remember just the green ones, the light green ones? Now it's been dotted with darker green. That's when they're talking. So the darker green ones, patches all over, means in that particular quarter, that airline used the word "capacity discipline." And so what we are going to do is suppose -- imagine -- imagine, as I was saying, two airlines -- two legacy airlines serving a particular market, we see if both of them were communicating or talking about capacity discipline, and what happens subsequently with respect to their capacity?

Just the summary statistics of how often the capacity discipline is used, the thing to notice, the first one is the legacy. So we have 253 quarters out of which they talk -- at least one of them talks about

19

capacity discipline 54 percent of the time. The LCCs, the local carriers like Southwest -- I have broken down Southwest here separately -- they don't use the word "capacity discipline" as often as the legacy carriers do. And, in fact, we don't find any effect of -- you know, we don't find any effect that, you know, the local carriers are reducing the capacity. Overall, the last row, about 38 percent of the time, airlines are using capacity discipline.

Now, the text data is out. This is the airline data. We have Bureau of Transportation -- so this is the T-100 domestic segment. As I said, we also augment that with the OAG market intelligence data, which gives us -- for every flight, we know who was operating for whom, and period of interest for us is 2002, the last quarter, until 2016, the last quarter.

The market definition, so it's a big deal to define a market, how you define a market, and so we go with the airport pairs. So basically if, you know, D.C. would have -- let's say if you consider D.C. to have three airports, but for us, a market would be Raleigh-Durham to BWI. That's a market.

So in the paper we also look at city pairs, so instead of thinking about each of these airports separately, we can consider, let's say, Raleigh-Durham

20

area and D.C. area. So we also have that result in the paper. Things become a little bit involved, you know, because there's inter-airport substitution, but for the talks today, I am going to just focus on airport pair.

So this is the construction of the variable of interest. This is what we will define to be communication. So we call that capacity discipline. So capacity discipline in market m, t is a product of two dummy variables. The first one is talk-eligible. By talk-eligible, we mean that there have to be at least two legacy carriers in the market for there to be even a question about communication, so we need the talk-eligible. That's where the market structure becomes important, and that's where the possible endogeneity will also come.

And the second part is all of these legacy carriers, at least two of them, there are at least two of them, all of them are using capacity discipline in the previous quarter. So if both of these are satisfied, then we take the dummy, and that's when we call there is a communication among these airlines, okay?

And so this is the basic regression model, very simple. On the left-hand side, we have log of seats

21

by Airline J in Market M in Period T. We regress on a bunch of variables, but the object of interest, the variable of interest is capacity discipline, which is the first one, the coefficient beta naught. We also control for -- you know, because -- you know, we allow for -- we control for various other factors that could influence the decision or the seat.

So, for example, we treat markets where at least there are two versus only one legacy carrier as separately, so that's the talk-eligible. Remember that the talk-eligible, which is the beta one, is also in the capacity discipline, okay? So think about the capacity discipline as an interaction between talk-eligible and everybody communicating.

We also control for monopoly. We don't know why the datas were missing, so we also control for the missing reports, those black dots in the picture, and we have a bunch of fixed effects. We have airline market fixed effects, airline year quota fixed effect, origin fixed effect, destination, time fixed effect.

So basically, what are we doing with the fixed effect? The intuition is that we are trying to control for any other demand-related shock that would affect the left-hand side variable, okay? And so our null hypothesis -- so we are interested in the -- in

22

the est beta naught (phonetic) and what is the sine of the beta naught.

So this is the -- so just look at the first column for the time being, and the object of interest is the first variable, and we see that whenever the airlines communicate, they reduce the capacity by 1.49 percent, and any statistically significance. And if you break -- you know, there are various ways in which you can decompose these effects. If we -- I am going to go focus on this one in view of the time.

So imagine -- so when we think about a market, a market is a mixed market if it's served by both the legacy carrier and the local carrier, and we decompose in the market what is the effect of communication on the capacity of -- you know, of the legacy carriers versus local carriers, and we see that, in fact, the local carriers do not have any effect. So the second one is, in fact, insignificant. The first one, the legacy, is -- becomes more stronger, okay?

So in a sense the summary of this is that there is some evidence that whenever airlines talk about capacity discipline, they do, indeed, reduce capacity, and it's only the legacy carriers who are doing that.

Possible concerns? As I said, financial transparencies, what we mean by that, just to repeat

23

what I just said in a bit, are carriers just being transparent with the investors about future plans? And so if this is, then there is really no reason for concern and this is not about capacity discipline helping airlines coordinate or collude.

The second is conditional exogeneity. There could be some unobserved factors that is affecting and driving our result, especially related with the way in which we talk about the text, okay?

And the third one, the third concern is the market structure being endogenous. The fact that a market is talk-eligible, meaning that at least two legacy carriers, or three or four, could itself be endogenous, which leads that the way in which we define the capacity discipline would be endogenous, and so we look at -- we address this by using a control function, okay?

So just a quick result of what we do. So we asked the following question: Do legacy carriers reduce capacity when they're the only carrier in the talk-eligible market talking? So if you are the only one -- suppose there are two airlines, two legacy carriers, and you are the only one talking, and you say, "I want to reduce capacity, I need to be disciplined," and this was a communication to the

24

investor, you would see a subsequent reduction, but we find none. So the first one -- sorry, the first row, only J talks, we find that, in fact, the effect is positive. So they don't reduce capacity.

Second, what about monopoly markets? So imagine a market where you're the only guy selling and offering the air service, and you say, "I want to reduce capacity," do you see any reduction? The first one -- the first row, you see that there is reduction. In fact, it's positive.

The third one we do is suppose -- as I said, we -- by definition, you know, if you look at the Awaya/Krishna paper and you repeat it again, the communication involved everybody talking, and suppose that only one person is left out. Suppose N minus 1 people talk, but 1 percent doesn't talk, what happens to -- what is the effect of that on capacity discipline?

If -- again, if it was all about transparency, then we should have seen some reduction, whatever the size might be, but we find, in fact, there's a positive increase. So this seems to suggest that there is something happening with the capacity discipline that needs more thought.

Okay, so just to summarize, we find that capacity -- that carriers do not reduce capacity when they unilaterally discuss capacity or when it's a monopoly market or if N -- you know, N minus one people or airlines in the market discuss the capacity discipline, but only one does not.

Conditional exogeneity, so this is a little bit involved. So suppose -- what are we worried about is that the way that we define and choose the word "capacity discipline," it could be that we're worried about that can there be other words that are positively correlated with the capacity discipline but negatively correlated with the log seats, okay, and that's what is basically driving it, because we are not controlling for that, and that's the big worry, because, of course, we don't know what exact words these guys are using. It could be that we're missing some other words, that it's left over, and it's not about capacity discipline. That's something that we would have to worry. And so to address that, we follow a test motivated by Hal White and Corrine Chalak. And so I am going to skip all of this.

Basically what we want to do, we want to -- we look at the text data, and we found -- and we look for any word that is related with capacity discipline semantically, and then we ask, suppose now you

introduce that word that is related to capacity discipline, and it occurs as frequently as capacity discipline, and you include that as an additional regressor.

And we would expect -- if our capacity discipline is capturing everything that we think it's capturing, then the coefficient on that should be non-negative, right, because that's what we worry about the most, and if it is not negative, it should be insignificant, right?

And so when we look at this, we -- this is -- this is what we find. I have one minute. So six words which satisfy -- they will use as frequently as "capacity discipline," but when you put it as an additional regressor, the coefficients are all positive, except for "slow." And it's a little bit big, and -- but the thing that is reassuring for us, at least, is that the capacity discipline in all of these regressions have similar coefficients, so it's kind of stable, even if you add and throw in all these variables, which might be semantically related.

And, okay, market structure, next, thinking about the fact that the talk-eligible is correlated. So we use -- I am going to go fast -- so we define hubs as -- hubs for each airline each month, and we

look at the distance from an airport to the hub, which is a proxy for the cost of entering, and use that distance -- so we predict what is the -- what is the likelihood of a market being talk-eligible, and then we redefine our communication.

I am going to go -- skip all these pictures. And so when we use the -- so basically just the control function, we find that the capacity discipline now with the control function is still significant, slightly smaller in size. So it's 1.14 instead of 1.45.

I think I just hit the button, so that's the conclusion. Thank you.

(Applause.)

MS. CARLSON: Next we will have Gloria Sheu, from the U.S. Department of Justice, Antitrust Division, to give a discussion.

MS. SHEU: Okay. Well, first, I want to thank the organizers for inviting me to do this discussion. I had a lot of -- I really enjoyed reading this paper. I thought it was really interesting on an important subject. I also have to start with the normal disclaimer, that the views I am going to express today are entirely mine and should not be purported to reflect those of the U.S. Department of Justice.

Okay. So as you just heard Gaurab discussing, I think this paper has two main parts. The first is the authors, they document an interesting empirical finding, which is that when you have two legacy carriers or actually all the legacy carriers in a market discussing capacity discipline in their earnings calls, you then see them actually decrease the number of seats in overlap markets. Then, the second part of the paper is to go on and argue that, based on that, this is evidence of collusion between those legacy carriers.

So given that that's the structure of the paper, I think the authors rightly go through a lot of work to try and say, hey, there's not an alternative explanation, like our story that this is collusion is what you should believe, and overall, I found that compelling. I thought it was particularly helpful to see cases where this pattern was not happening. So, for example, low-cost carriers did not appear to be participating in the markets where not everybody was discussing this, as Gaurab was just saying. That's not -- does not appear to be affected. So that kind of, like, maps out, like, exactly what it is we're talking about here, but nonetheless, it's, of course, difficult to prove a negative. I think that's typical of this type of paper.

As we saw, there was a lot of kind of different steps and additional work that the authors did to try and cross off as many of the alternative explanations as possible. Today, I think rather than add to the stack of things that they've already tried and I'm sure that they're considering trying, I want to step back a little bit and discuss a little more, like, the wider antitrust context for this type of research.

So one reason why I really like this paper is I think it's super important for people to work on research related to collusion and coordinated effects of mergers. On the one hand, we see that I think antitrust practitioners -- so folks who work as experts in, for example, merger cases -- and courts, finders of facts, have converged or somewhat converged to a generally accepted set of ways of thinking about the unilateral effects of mergers.

We see, for example, similar types of merger simulations come up in a lot of cases, and that's great because it means when you're looking at a new case, you have some common ground to think about. You're not starting from zero with your analysis.

On the flipside, for coordinated effects of mergers and collusion and non-merger cases, I think

it's way more wide open. There's not quite as, you know, a generally accepted set of models or empirical tools, and it's not that these cases don't get brought and that people don't look at these things. That definitely happens, but at least from my experience, the cases that I've seen, a lot of the work ends up being very specific to that particular matter, and it's -- and then the work is kind of, like, difficult to transport into another situation, which then, if you're looking at some other potential matter, you're kind of starting from zero, and you might not be on the same page as other people who are thinking about the same investigation.

So this is largely an empirical paper. Gaurab referred to some of the existing literature that -- theory literature that could underpin it, but I think it's helpful to think about that some more. I think some -- in the paper as it's written now, I think maybe some more, like, light exposition along those lines -- that would be my one comment -- would help fix ideas on this a bit more, but I am not suggesting some sort of separate theory or extrastructural estimation, as that would be an entire paper unto itself.

But what I found really interesting about the

collusion that they're identifying here is that it's kind of partial, and that could be interesting to rationalize in a model, and there are models out there that deal with this, but in this particular instance, what we have to have is a model that would have only certain geographic markets being affected at any particular time, so specific overlap markets where they're communicating. Some firms are not participating in all time periods, so some of them are not saying "capacity discipline" at all times. And some firms appear to be entirely excluded, so the low-cost carriers are not involved at all.

So you need a model where somebody or a group of firms find it in their interest to collude, but they don't want to, like, collude too much, right? They don't want to do it all the time, for all time periods and for all market conditions, so there's got to be some sort of friction in there.

So just breaking that down a little bit more, for the geographic markets, like, why were these particular markets chosen would be an interesting question, and I would think that in the model that would underpin this, you might find that there was some sort of incentive compatibility constraint that was -- meaning that certain markets got excluded.

The thought process here would be, all right, some -- maybe they would have wanted other potential overlap markets to be included, but for whatever reason they didn't all communicate, and maybe that was because maybe they didn't want to include it or maybe it was actually in -- somebody would have wanted to, but if they had done that, they would have induced cheating. That would be one thought process that could rationalize that.

Stepping even farther back, like, generally, what kind of punishment would you set up to get this? Some measures of capacity are publicly available. The data was used in this paper. I think that these airlines are reasonably well informed about what's going on around them, so monitoring might not be a huge issue. So why at some points aren't they talking about capacity discipline? Is it a situation where, you know, they all just decided we're not going to do this right now, or is it that the scheme broke down and they went into a punishment phase. If it's the latter, how did they start up again, right? So that would all have to be built into the structure here.

And then another really interesting thing is we had these LCCs hanging out here. You see in the antitrust literature, it talks about mavericks that don't participate in a collusion scheme. Could these have been mavericks? It's possible. That would be a situation where the legacy carriers would have wanted them to participate but couldn't get them to for whatever reason, and the LCCs might have actually prevented additional collusion or the collusion that happened from being more successful.

The flipside could be that these LCCs just weren't that good substitutes or weren't that high a competitive constraint for these legacy carriers, so they just didn't bother dealing with them. I could imagine that the answer would depend on the market and maybe on the LCC they were talking about. So there could be some variation there that might be interesting. And I think that, you know, the idea of a maverick is something that pops up a lot in antitrust contexts, but trying to actually, like, look at something empirically on that actually might be really helpful.

And, of course, there's just the wider, really big-picture, general questions. Any time we're looking at antitrust relevant research, there's a question of did mergers play a role. We definitely had a bunch of airline mergers in the not-too-distant past. I mean, this could be some additional empirical

work, and the paper just looked specifically at what happened around those. Did the firms involved change? Did the markets involved change? Did the amount of talk change? Did this bring certain firms into the fold? Honestly, just some empirics around that itself might be interesting, and then that could also say, okay, maybe that would be something that would be interesting to model.

And then, of course, the million dollar question is, what happened to prices and consumer welfare? Again, that would be something that you'd need some structural modeling and estimation to get to, but that is really the question that we're after when we're thinking about looking at collusion and, for example, making a case as to why something might be prohibited conduct in a court of law.

(Applause.)

MS. CARLSON: So I'll ask Gaurab to come back up to the podium. We will have about ten minutes for questions. Alex and Jenn in the back there are wandering around with microphones. So if you would like to ask a question, just flag one of them and we'll have some discussion.

MR. BRUESTLE: Steven Bruestle, Federal Maritime Commission. This reminds me of a growing

antitrust literature on the effect of the same investor owning shares in multiple competitors. Investors like to diversify their portfolios, yet this can encourage collusion.

Do you have a way of controlling for this? Perhaps overlapping owners impact the coordinated effect that you found.

MR. ARYAL: We did look at the timing, meaning who initiates the questions that leads to the answer that contains capacity discipline in the hope that we could probably try and tie the analyst, let's say, back to the real owners, and we didn't find any effect of that at all, but we did not pursue seriously, to be honest, the line of common ownership yet. But we are aware that that's something that's happening and, yeah, we don't know really.

MR. BRUESTLE: Okay, fair enough.

MR. RASMUSEN: Hi. I'm Eric Rasmusen, Indiana University. I wonder if you could tell us more about earnings calls. Do they happen every quarter? And what order do they occur in? Is it the same airline who goes first? That sort of thing would be really interesting to know about.

MR. ARYAL: Right. So it happens every quarter, and as I said, we did try to look at the

timing issue. It didn't matter if Delta was the first to do the earnings call in that particular quarter. Is that what you mean?

MR. RASMUSEN: Oh, yes. Oh, and also, do analysts ever bring up capacity discipline?

MR. ARYAL: So that's kind of related to what I was just saying, that analysts do bring up sometimes the capacity discipline, and we do try to look at if there was any -- you know, if we could see some pattern, but we did not find any pattern, either when the analysts bring it up or one of the legacy carriers is the first one to bring it up.

The idea that we had was that there's this new paper that is coming in AER where they show, in Australia, that BP kind of leads others to collude and follow them through, and we did try to see if there was any evidence of any particular airlines doing that, but we didn't find any.

And to be honest, I guess, the idea of cheap-talk and communication with a leader is also -- we don't know how to conceptualize that idea. So, theoretically, we don't know what a model -- you know, can there be, you know, a theory model where a leader leads others through the cheap-talk, you know, where monitoring might be a little bit messy or at least

with the lag, so -- but yeah, thanks.
AUDIENCE MEMBER: Hello. So I guess the main concern I guess is that when these four -- when all the legacy carriers talk, you know, it's a big negative demand shock. I guess that's one way of thinking about this. And then one way of reading the results is that those markets where at least two carriers are present are maybe more cyclical, more affected by those aggregate demand shock, and, you know, that -- I mean, you know, if, for instance, you figure Philadelphia is probably less cyclical than LAX, LaGuardia, and that might be what's going on, and one way of accounting for this would be to add more controls, I mean, for those things, and I'm wondering if you --
MR. ARYAL: So, yeah, add more local controls or --
AUDIENCE MEMBER: Well, I mean, you could easily add those aggregate demand shock controls.
MR. ARYAL: Sure. So at least the first part, we do -- because, you know, we also control for the talk-eligible, so that takes care of maybe the first part of your point, that really we are interested in the interaction of who's serving and if they are all talking versus who is serving, so that takes care of

that point, but we did not add any global -- you know, maybe we could use past -- something like, you know, past demand or past load factor or something like that to get at.
We did have a time trend at the -- at the airline and the market level, so that also takes care of if -- you know, this is -- the thing -- the way we are thinking about it is is that if the demand is growing at 3 percent, does the capacity also grow at 3 percent or not when they're talking? I think that's how we interpret that. So it's -- so we have to think about exactly how the identification would work, but that's something that we could think through. Thanks.
MS. FORBES: Hi. Silke Forbes, Tufts University.
I was wondering if you could talk a bit about market definition. You said -- you talked about the results using airport pairs.
MR. ARYAL: That's correct.
MS. FORBES: What happens when you use city pairs instead?
MR. ARYAL: So in city pairs -- so what matters is whether you have two or three airports make a difference, and so, you know, we tried to look at why is there a difference between two and three. So, for

example, if you have three, then there is no effect, but if there are two, then we find that there is an effect. The closest that we could come up with was the hoteling thing, where, you know, things are a little bit less stable when you have three. I don't know, that's just a -- I'm stretching here, but we -- yeah, with three, something happens, and we don't find any effect. You know, we do that in the paper, so it's -- thanks.
MR. LAU: Hi. My name is Yan Lau from the FTC. I just noticed that you might have tried this before, but have you tried -- like, you have a log specification, but if you were just to go with a linear specification, you can actually put in all the zeros of the dependent variable, and then you could potentially get at market entry and exit, because right now I think what you're doing is you're throwing away all the routes where an airline has zero seats, zero capacity, and so if you put in all the zeros and you're willing to get away from the log/linear specification, then you can see whether people exit or enter the market based on capacity discipline.
MR. ARYAL: Okay. We did not -- we did not do -- we never thought about linear specifically at all, to be honest, but I have to think -- so, yeah, I

have to think about -- I don't know. I don't want to say anything without thinking about it, but thanks. I'll think about it.
MR. LEWIS: Eric Lewis at DOJ. So thinking about Gloria's comments, I think in a model of collusion, we think about kind of three states of the -- three possible outcomes, which is either the baseline or collude or punishment, and so you really just have those two, either you're cooperating or you're not, and your control group is sort of nesting in possible baseline states with also cases where there might be punishment.
And so I wonder if that is exacerbating the difference, or I guess the bigger question is, is thinking about -- you know, given that this is a repeated interaction, it seems like you could probably do something a little bit richer to think about how the outcomes depend not only on what -- just the communication in the previous period but also what were the previous periods' outcomes?
MR. ARYAL: So I don't know the answer to the first one, I have to think about it, but for the second one, we did -- at one point, we did try to redefine communication as a continuation, sort of like, you know, without any break, if everybody's

serving the market, everybody's talking, and at the time we -- the effect was actually much stronger.

But, again, the difficulty is -- I think the difficulty that we had conceptually was to map it to a model, and as Gloria said, there isn't any model that really fits the market, so we had to make a choice, and we decided to just play it safe and say less than what we possibly could if we stretched the market a little bit. But your point is well taken, yeah. We could do something much richer and cut the data in many different ways, which we have not done, and I think -- but before all of that, before we do all that, if we have any energy left, we would probably devote it to think about prices. I think that's probably more important than anything else, and we haven't done that. We don't know how to do it, to be honest.

MR. SINGER: David Singer, Northwestern.

I could imagine -- sorry. I could imagine a model where, let's say, demand is really price in a market, is really price-inelastic, and the firms are up against their capacity constraints, and just a little bit of cut-back in capacity would be really good for each firm, and that the cutting back capacity really could be a noncooperative equilibrium as

opposed to a collusive equilibrium.

So I'm wondering, is there any way in your data that you could begin to get at when -- begin to get at the idea that the cut-back could only arise out of collusion versus, you know, as just part of a Nash equilibrium in capacities noncooperatively.

MR. ARYAL: Great question. We did try to look at how these estimates change when we -- separate markets by business passengers, with the idea that business passengers have lower elasticity, and so if your market has a higher fraction of business travelers, then things would -- and -- but we did not really use that to think about that the reduction in capacity could arise only out of collusion. We did not -- we did not really -- but that's a good -- a really good point. So we should. We should.

MS. CARLSON: Okay, I think we are out of time (off mic). Thank you.

(Applause.)

MS. CARLSON: Great. That was an excellent discussion.

(End of session.)

PAPER SESSION:
ONLINE PRIVACY AND INFORMATION
DISCLOSURE BY CONSUMERS
- - - - - -

MS. CARLSON: Our next paper will be presented by Shota Ichihashi from Bank of Canada. He will be presenting Online Privacy and Information Disclosure by Consumers.

MR. ICHIHASHI: Thank you very much for having me in this great conference. About myself, I am Shota Ichihashi. I finished Ph.D. this summer, and I am doing microeconomic theory in the Bank of Canada. Today I'm talking about online privacy and information disclosure by consumers, where I asked the following question:

What are the welfare and price implications of a consumer's privacy in online marketplaces? There are many ways to tackle this question, but the following is what this paper cares about.

There are online sellers who observe detailed information about consumers, say browsing, purchases, or their characteristics, but consumers can potentially affect to what extent this information is revealed. For example, they may delete cookies to hide their web browsing activities, or if they are

more sophisticated, they might create multiple accounts on shopping websites to obfuscate their purchasing behavior.

So what I want to capture is the interaction between the consumer's incentive to reveal information and the seller's incentive regarding how to use the information, and to that end, I consider a simple model. So this is a theory paper, a simple model, where a consumer discloses information to a seller who uses the information to make a product recommendation.

Now, what's that tradeoff? So there are many reasons that consumers may or may not want to reveal information. There can be some intrinsic privacy concern, but this is not what this paper is about. So what I study in this paper is the following economic tradeoff.

The benefit for the consumer, the benefit of disclosing information is that the seller can learn about their preferences and recommend or advertise more appropriate products, and the downside is, as you might expect, is a potential price discrimination. Sellers may base prices on what they learn about consumers who capture more of the surplus. And I will show how this tradeoff shapes the consumer's incentive to reveal information and the seller's incentive regarding how to use it.

Now, I don't have much time to cover the detailed literature, but let me just say this is the intersection of the theory literature, information design, down right, and the literature of the economics of privacy, the rest of that box. But I'm happy to talk about the marginal contribution offline.

So it's a theory paper, so I will show you a model, and I show many results, and this is the first half of the paper. Later, if time allows, I'll consider -- I'll talk a bit about the second half of the paper as an extension.

So let's begin with the baseline model, which is quite simple. There is a single seller and a single consumer, but the seller sells two products, product one and two, and the consumer eventually buys one of the two products or nothing. And u1 and u2 are the consumer's variations for products one and two, and they are IID, nonnegative and nondegenerate, and there is no production cost.

Preference is standard if consumer buys product k, his payoff is value minus price. If he buys nothing, he gets an outside option of zero payoff. The seller's payoff is its revenue, and both of them, both players, are risk-neutral.

Now, so, this is the primitive of the model. Now, let's see the timeline of the game. So I will show timeline, but what I want you to remember from this slide is I consider two models, two different games, that differ in whether the firm can price-discriminate.

In one model, a model of nondiscriminative pricing, the seller sets a price for each product at the very beginning. And then taking prices as given, the consumer discloses information. The seller learns something about his preferences, and then the seller makes product recommendation. Consumer makes a purchasing decision, and the game ends.

The other model is a model of discriminatory pricing. The only difference is that the consumer discloses information and then the seller sets prices. So we consider two models that differ in the timing at which the seller sets the prices.

Now, if this is clear, let me talk about three components which are common between the two pricing regimes; namely, disclosure, recommendation, and purchasing decision. So let me begin with the consumer's information disclosure.

So this slide summarizes information disclosure and is probably most important slide in the setup. So

in the model, the math is a bit broken, but before observing the product in u1 and u2, the consumer chooses a disclosure level, which is number delta between half and one, and then after this choice, the seller observes delta and a signal realization.

Signal realization is a random variable whose distribution depends on which product is more variable to the consumer and this delta itself. So, namely, if this diagram is correct, whenever product one has a higher value with probability delta, signal one is realized. Whenever product two has a higher value with probability delta, signal two is realized.

So one important observation is that the greater disclosure level delta the consumer chooses, the more accurately the seller can learn about which product is likely to be the best. So intuitively -- so delta is a precision. Delta is how much personal data the consumer discloses, and implicit assumption is that if the consumer disclose more, the seller can learn more about which product is more likely to be suitable to the consumer.

So what's the interpretation of signal one and signal two? So signal one and two indicates the consumer is more likely to love one product than the other, and how this particular signal looks like

highly depends on the particular application. I don't cover the application in this talk, but basically this diagram summarizes a situation in which the consumer can affect how precisely the seller can learn about his preferences.

And, again, let me just emphasize, when the consumer chooses delta, he doesn't know his valuations, so we don't need to worry about cheap-talk problem or we don't need to worry about the possibility that consumer manipulates signal one and two ex post.

All right. So this is information disclosure. Consumer chooses Delta. Seller learns which product is more likely to be best with a different precision. And then, after that, again, regardless of the pricing regime, the seller updates his belief and recommends one product, and the consumer sees the varying price of the recommended product, and he decides whether or not to purchase the recommended product.

Recommendation can lead to relatively which product to recommend, or depending on application, it can be which online ad to display, as I discuss in the paper, but in terms of the formulation, what's important here is that while the seller sells two products, it can only recommend one product, and the consumer can -- a consumer decides whether to buy the recommended product -- namely, he cannot purchase the product that is not recommended -- and so this is not the result of seller's optimization, of course. This is an assumption.

With this assumption, consumer can only evaluate the one product, and what this tries to capture is the situation where the consumer's attention is limited; namely, compared to the variety of the whole products, here two, the consumer can evaluate only a small subset of the products, here only one. This particular formulation of limited attention is in line with the -- some theory -- decision theory approach of limited attention. One twist here is that it is the seller who affects what products the consumer pays attention to.

All right. So this is basically the whole timing of the game. So let me just wrap up the setup. Under nondiscriminatory pricing, pricing comes first, product one and two, and then consumer reviews the information delta. Then seller learns about which product is more likely to have a higher value, so recommends product. Consumer sees the value and decides whether or not to buy.

Under discriminatory pricing, after information

disclosure, the seller decides which product to recommend with what price, and I consider a subgame profit equilibrium with some tie-breaking rule. Now, if this is clear, let me move on to solving the model.

So I will solve the game backward. So I will show that how recommendation and pricing look like, and then I will show the entire equilibrium. This slide summarizes the seller's equilibrium recommendation strategy, which is quite intuitive. For example, given any delta, signal two suggests the consumer is more likely to have a higher value for product two. So regardless of which pricing regime we focus on, the seller optimally recommends product two.

And from the consumer's perspective, this means the benefit of more information disclosure. The higher delta disclosure level implies he's more likely to be recommended the best product. If the delta is one, he surely recommended whichever product has a higher value.

The question is, how does this more disclosure and better recommendation affect product prices, which is -- to the consumer, this is relevant under discriminatory pricing. So suppose the consumer increases delta from some number above half, which means the seller is more likely to recommend the best

product, the product with the higher value. The value is maximum of u1 and u2, and how does this affect the consumer's value distribution for the recommended product, and how does this affect the pricing?

Now, I don't show the math, but basically if the consumer discloses more information and if the seller is more likely to recommend the best product, consumer's value distribution for the recommended product has a lower hazard rate just using the property of the higher and lowered ordered statistics of two random variables, and this intuitively captures the idea that the consumer's demand for the recommended product becomes less elastic. Therefore, the seller, monopolistic seller, charges a higher price.

So what happens in this model under discriminatory pricing is that if the consumer discloses more information, then recommendation gets better, which in turn implies the consumer demand becomes less elastic, which gives the seller an incentive to raise a high price.

Now, so, what we've seen is more information disclosure leads to the higher price -- no, more information disclosure leads to better recommendation first, but under discriminatory pricing, it leads to

the higher price, and by combining these observations, we get the first result. Each pricing regime has a unique equilibrium, which I show in the paper, and the seller is better off and the consumer is worse off under nondiscriminatory pricing.

In other words, the seller prefers to commit not to price-discriminate, which makes the consumer worse off. This is a little bit different from what we might imagine it from the standard price discrimination model, so let me give intuition.

So as we saw before, under nondiscriminatory pricing, a consumer chooses a disclosure level, taking prices as given. So what he cares about is just a recommendation accuracy. So it is optimal for the consumer to set the highest disclosure level to make sure that he's recommended the best product.

However, when the seller sets a price for each product up front, the thing is like this. When the seller considers what price to set product one, because the seller knows consumer is going to disclose much information with which the seller can make accurate recommendation, so that this product one is going to be recommended only to consumers who have high valuation for it. So, therefore, the seller sets a relatively high price for each product, one and two, in advance.

Now, in contrast, when the consumer discloses the information first under discriminatory pricing, the consumer is in some sense the first mover (indiscernible) leader, who can chose a disclosure level, balancing the benefit from better recommendation and a cost from a higher price. As a result, he chooses a weakly lower disclosure level by which he can enjoy a weakly lower price and a higher payoff, although recommendation can be a bit noisier.

So what happens is the seller wants to commit to nondiscriminatory pricing, which encourages information disclosure, but in this pricing regime, consumers disclosing too much information in the sense that if the consumer could precommit to withhold some information, he could be better off.

Now, this intuition is based on the relative commitment power or the timing of moves between the seller and the consumer, but today I'd like to show that another intuition, which is based on the following alternative interpretation of the model. So only in this slide let's forget about the single consumer model, but imagine there are a unit mass continuum of consumers, and in this interpretation, in this formulation, under discriminatory pricing, the

seller can charge different prices to different consumers, and, of course, recommend different products to different consumers.

Now, under this nondiscriminatory pricing, still the consumer disclose information first. Many consumers disclose information, and then the seller sets a single price for each product and then make recommendation. So there is a slight difference. There are many consumers, and difference of pricing regime is whether the seller can personalize prices.

In the paper, I argue this is essentially the same as the original model we've seen, and, in particular, consumers are worse off under nondiscriminatory pricing, but in the current alternative interpretation, we can think of this as a classic tragedy of the commons due to the following negative externality associated with information-sharing. So here's what I mean by negative externality.

So suppose there are 1000 consumers, nondiscriminatory pricing, and suppose there's 100 consumers disclose more information. This gives the seller an incentive to charge a higher price for each product because, on average, the seller can recommend the better product, can give a better recommendation

to a larger fraction of consumers, but these higher prices lower the welfare of other 900 consumers who might not decide to disclose information. So what happens here is information disclosure by consumers lower the welfare of other consumers through higher prices when the seller cannot personalize prices. So in equilibrium, consumers disclose more information than what would maximize the joint surplus.

So in contrast, if the seller can personalize prices, each consumer take into account the impact of his disclosure on prices, so in total, consumers disclose weekly less information, and they are collectively better off. So this is a bit outside based on the alternative formulation. So let's get back to the original single consumer model, and this is the same slide as the previous, previous slide. So let me give you two relatively straightforward implications of this result.

So, one, this gives a seller a rationale for committing not to price-discriminate. Of course, there can be many, but one story is that once the seller starts to price-discriminate, consumers are discouraged from providing information, and this lowers the matched quality between the products and the consumers and hurts revenue.

As this intuition suggests, it is really important that there are multiple products, important that the consumer cannot evaluate all products, and also important that consumer can affect how much information to disclose. So this highlights the key variable here is the fact that consumer can affect how precisely the seller can learn about themselves, learn about consumers' preferences.

The second is a little bit on policy side, so I don't have a specific regulation in mind, but the observation that consumers disclose more information than what would maximize their joint surplus suggests there might be some regulation which limits the consumers' disclosure or the regulation which limits the amount of information that sellers can seek for which benefits consumers. So this is because such a regulation might restore the consumer's commitment power to withhold information from sellers. Now, so, these are the two implications, and we cover the first half of the paper.

Now, let me spend the rest of the time to talk about the second half of the paper, which -- where I study the following unrestricted model, which is the unrestricted version of the model. So, first, the seller doesn't just sell two products. Seller can

57

sell K products with IID values, and more importantly, the consumer can disclose any information about this K dimension random variable or vector. So I don't show you a formulation of what I mean by "any information," but in this model, the consumer can, for example, disclose the name of the -- consumer can let the seller learn which product has the lowest value, or the consumer can let the seller learn his willingness to pay for a particular subset of the products. So the consumer's choice set about disclosure is extremely rich.

Now, why do I consider such a situation? Absolutely not because I think this is the most realistic, but because I want to see, one, robustness check of the main finding with respect to the assumption of what information the consumer can disclose. Here, consumer doesn't just disclose the information of which product is better, but he can potentially let the seller learn his vertical willingness to pay.

And another important reason is that this connects the paper to the theory literature. In particular, the recent AER paper by Bergemann and Brooks and Morris, 2015, which is basically the -- so my single product version model, this model, feeds

58

their work.

Unfortunately, I don't have time to analyze the model, but the punchline is we get the same result. Whenever product two -- no, whenever the seller sells multiple products, the seller is, again, better off and the consumer is worse off under nondiscriminatory pricing, and also under a very mild assumption on the distribution -- the variation distribution, we can conclude seller is strictly better off and a consumer is strictly worse off under nondiscriminatory pricing.

So the proof is much, much longer, but basically I show that nondiscriminatory pricing, as we can expect, it has a benefit of encouraging disclosure which leads to better recommendation, but to the seller, there is an obvious loss, which is the seller cannot tailor prices on information. So what proof shows is the benefit dominates a loss from the seller's perspective.

But actually, in the paper, I can never derive the -- what information the consumer reveal under discriminatory pricing, so without knowing the disclosure policy, I compare the seller and the consumer welfare, and in the middle step, I characterize the most efficient disclosure policy of the consumer.

59

And the result theorem is in contrast to the Bergemann, Brooks, and Morris, in the sense that -- in the sense that in this result, the seller has a strong preference toward nondiscriminatory pricing, under a mild condition, but in the single product version, the seller is indifferent between two pricing regimes, which I think itself is interesting.

Now, so, the -- what's left, I haven't talked about the sum of consumer and the seller total welfare. One natural question is can nondiscriminatory pricing, which the seller prefers, enhance total surplus? The answer is no if there is only one product. The answer is it depends if there are multiple products.

In our paper, I formally show nondiscriminatory pricing always leads to the more efficient recommendation, never -- there is never product match -- product mismatch under discriminatory pricing. It always leads to the highest probability of a trade. So it depends on which effect dominates, which pricing regime is more efficient. If there are many products, then the first effect dominates. If there are many products with IID values, eventually nondiscriminatory pricing leads to a greater total surplus because it encourages disclosure and leads to

60

better recommendation.

All right. I have more than one minute, so let me cover. So another question of our thinking is can there be some institution which can improve the total welfare farther, and one thing we often talk about is when it comes to personal data, the rough idea is if there is a market for data, it can enhance the welfare. There is a -- this is a naive way of incorporating market for data in my model.

So basically, in addition to the model, I explained at the very beginning the seller can offer the consumer to purchase information and make a -- pay money. What I show is that under nondiscriminatory pricing, this additional stage of buying data has no impacts, because consumer's happy to disclose information for free.

Under discriminatory pricing, this additional stage of asking information and pay money can improve the seller's revenue without affecting -- without lowering consumer's payoff. So in this case, seller asks the consumer to reveal full information and to make a transfer, which keeps the consumer just indifferent between accepting and rejecting the offer, and as a result, there is a perfect price discrimination, which is efficient.

Now, let me wrap up. So the question I'm interested in is what are the welfare and price implications of consumers' privacy, and the conclusion, under some assumptions, seller is willing to commit not to price-discriminate, which hurts the consumer but may improve total welfare. Thank you so much.

(Applause.)

MS. CARLSON: Thank you.

Next we will have Guy Arie from the University of Rochester, Simon School, to give a discussion.

MR. ARIE: All right. Yeah, so thank you for the organizers. It's a wonderful conference. And thanks, David, for inviting me. And, Shota, it's a very interesting paper.

So very quickly, so what is this paper about, right? It's how do sellers use the buyer's information. And let me try to, like, give -- we're going to get to an example, but basically there's two things that he's trying to talk about. One is sellers can use the information to just offer you a better matching product, okay? And the second is they can use the information to price-discriminate.

And the point -- right, so, for example, I need to -- I get to decide what Amazon sees, like what I --

which news I see and whatever -- you know, Amazon knows a lot about me, and they can use that to decide which Halloween costumes to suggest to me, right, or they can use that to decide also how much to price the various Halloween costumes that they suggest to me, right? And so that's the two dimensions that he's going to talk about.

And the main point of the paper is I'd actually benefit not only from letting -- so we're not going to talk about do I benefit exactly from letting Amazon know everything about me, but in the world that Amazon knows enough about me, I actually benefit from them price-discriminating. So let's see how, and you can only see the -- something that someone like that -- it's about that.

So not going into the model components too much, just think about it -- you know, the way you really want to think about it is there's a monopolist, let's say that's me, and I am selling you two products, okay? Now, these two products have downward sloping demand curves, but it happens to be that half of you like this product more than you like that product, all right? So there's two downward sloping demand curves.

For half of the population, this one is higher.

For half of the population, this one is a little bit higher. The problem is, I don't know which one, okay? So that's how I start -- that's how we start this model. I don't know which one, okay?

So what happens if I don't know anything, okay? The problem is you can't go to the one you like more. That's the problem this monopolist is facing. You have to go to someone you don't know which product you like more, and I don't know which one to offer you, so I just basically can -- because IID, but I might as well sell everyone this one, okay, because I don't know, and as a result, my demand curve isn't that great, okay, because it's coming from the aggregate population.

So what happens if I know, okay, so full disclosure without price discrimination, so that's the first thing -- that's basically the baseline for this model, all right? So what he's saying is, okay, now what's going to happen is relative to the world that I don't know, so I'm facing -- as a monopolist, I'm facing this one downward sloping demand curve. And now I actually am going to know, so you're going to come, and I'm going to tell you you go right, and you're going to come, and I'm going to tell you you go left, okay, and actually have information to base that

on, okay?

So if I have information to base that recommendation on, what's going to happen is I can say, you know what, I have a -- I'm a monopolist. I have a downward sloping demand curve for this product. I have a downward sloping demand curve for that product. I'm just going to price it accordingly, you know, I don't care about the -- you guys not knowing, because you are going to know, okay? So that's the first result that we have, okay?

As a result of this, as a monopolist, I am doing better, right? The demand curve is higher, so I am getting more revenue, and, in fact, in this model, in this paper, I am going to get the best -- the highest revenue that I can as a monopolist, okay?

So if we think about this, there's two forms of price discrimination that are going to come in. Here, the price discrimination is completely independent both of the fact that -- how honest you chose to be with me as the monopolist and what you actually told me, okay?

In practice -- and I could, you know, price-discriminate based only on how honest you choose to be with me, and I can also choose to price-discriminate based on what exactly you told me. Because I don't discriminate on anything, I have two products, it ends up being in your best interest as a buyer to just tell me the best thing, because I'm just going to tell you where to go, okay? So all this centers on the line, and I am just going to be a very profitable monopoly, all right?

Notice that we don't need to have any information disclosure, right? What we could have is just there's a product here, there's a product there, they have prices. You, customers, go choose whichever one you like, okay? I am going to set the same prices, okay? So, like, fixing a problem here by full disclosure that we don't necessarily have to have if we don't have -- if we have enough information, and, of course there could be reasons that we don't have the information.

Another observation about this is we tend to assume, you know, before we start thinking about value rationality very carefully and all of that, we generally tended to assume that the customers are pretty well informed, right? So well informed customers tended to be the best, and here it's actually the worst case scenario in a sense. Other observations are less important for this discussion.

So this is the baseline, and now we can talk

about price discrimination, all right? So, in fact, I can price-discriminate, like I said, based on two things, how honest you are with me, so do you let me see everything you do, okay, at home? Do you have an Alexa, right? Or what exactly did you tell me, right? Did you tell Alexa that you really like wearing very scary costumes? So two different things that I can price-discriminate based on.

So if I only discriminate on policy, okay, what's going to happen? Well, you already know as customers that if I only discriminate -- if I don't discriminate -- if I don't discriminate at all, okay, then I'm basically going to act like a very, you know, vicious monopolist, right? So it's actually in your best interest to make me be scared a little bit, okay, to make me less certain as a monopolist.

So if I only disclose based on policy, some of you are going to have Alexa and some of you are not going to have Alexa, right, and I'm not going to know which is which. I only know how many Alexas there are in the world, and that's going to make me, as a monopolist, a little bit softer, because I need to handle the fact that some people are going to the wrong product, right, and that's going to actually decrease the price, okay.

So that's basically what happens if I only discriminate based on price, right, or that's the intuition. Then, if I discriminate -- what happens if I discriminate based on price? Well, if I discriminate based on price, given that I'm pricing as a monopoly already, okay, it ends up that you as buyers could do a little bit better, because you're not going to tell me, hey, price even higher, right? You are not going to give me -- as long as you have control for what you are disclosing, you are not going to disclose information that tells me to price even higher, but because you have control over what you're disclosing, maybe you could figure out a way to tell me, you know what, price a little bit lower for me, okay? Not for everyone, but for me, price a little bit lower, okay?

If you can commit initially to do that in a credible way, I'm going to go along with it, right? And that's the second main result of this paper, is that if there is full information disclosure, so if you have complete flexibility in how much information you disclose as buyers, then you are actually going to use that, if you could, to give me exactly the right information to make sure that no one goes away empty-handed, okay? As a monopolist, I am going to

actually give everyone discounts, that if I can't -- if I need to, to make sure that everyone buys something.

So now that we understand -- so, like, this is what's going on, what can we say about this? Well, this works because you, the buyers, have a whole lot of control here, right? You, the buyers -- it's not like you tell me, hey, here's everything about me. Alexa is in my house. You know everything about me, right? No, that's not full disclosure, right? Full disclosure here is here is exactly the right things I want Alexa -- like, I'm committing ahead of time. This is -- and I'm fooling Alexa, right, but I'm committing ahead of time to when exactly is Alexa going to hear things in my kitchen, okay?

And when I commit exactly, then it's only, you know, when my kids cry and I say, "You are not going to get a Halloween costume," and if they hear that enough, maybe they're going to give me a discount, okay? So in that world, okay, I can get some information disclosure in letting Amazon price-discriminate based on the information that they hear from me actually works, because I'm giving out the signals that say, hey, Amazon, I deserve a discount, okay?

But in the world where full disclosure is -- where Alexa hears everything, we might actually be closer to, you know, the firm knows exactly how much -- Amazon knows exactly how much I'm going to pay for a costume, and they are going to tell me the only costumes we have left are like, you know, 35 bucks. So what's going to happen? We're going to have increased efficiency, because Amazon is going to know exactly, for every one of us, how much we are willing to pay. So everyone is going to buy, okay? But, of course, you know, not a whole lot of surplus for consumers.

So the policy relevance, there's a lot of policy -- you know, there's a lot of things. The main thing I want to -- you know, in terms of, you know, FTC, competition, right? There's this Corts paper that I really like. Take all of this, throw competition in it, a lot of problems get solved, and even are better for customers, all right? So a lot of the results here are coming from the fact that we're starting from a monopolist. A lot of other policy things that we're going to skip because I wanted to really get across the intuition of the paper.

It's a really interesting paper of sales accounts. Really, it's a terrific piece of research

on really a difficult and important question. We need to think about this more, mainly on can we have models where the baseline model isn't this dystopian monopoly, all right? If we can stop for someone there, then we can see what happens with discrimination and so on, right?

That's it. Thank you very much. Great paper.

(Applause.)

MS. CARLSON: All right. So, I'll ask Shota to come back up, and we'll open the floor for questions.

AUDIENCE MEMBER: Yeah, we're just wondering if you can actually highlight the difference between your paper and the Roesler and Szentes papers in which the buyers can actually learn about the evaluations, or in this way I'm thinking that the signal structures that the consumers are committing themselves to is the same as, you know, learning about their own evaluation, since they don't know it anyways.

MR. ICHIHASHI: So that's the Buyer Optimal Learning paper.

AUDIENCE MEMBER: Yeah. Yeah, that one.

MR. ICHIHASHI: Yeah. So I think one way to think of this model can be a multiproduct version of their model. So in my model, the way in which consumer gets to the best product is to let the seller

learn about consumer's preferences, and as a revenue-maximizing seller, it recommends the best product.

And I naively extend their work, buyer learning, to multiproduct, I can imagine a situation where I commit to what I learn about the values of the two products, and based on the information, I go to whichever product gives me a higher profit. But like I said, if we think of the two products, multi version of the paper Buyer Learning, there may be a similar force in the sense that buyer wants to commit to learn less about the variations of the products by which he gets a noisier product match, but monopolist sets a lower price. Yeah.

AUDIENCE MEMBER: Yeah.

AUDIENCE MEMBER: So I'm wondering, so, for example, for the discussion -- the example was Amazon, and so Amazon is a platform that sells products that are made by other, let's say, manufacturers. So in that case, how would the model change once you take into account that the products themselves are actually produced by different manufacturers and they actually compete -- I guess what I'm trying to ask as well is that, how will I learn the incentives of the upstream in the platforms?

MR. ICHIHASHI: I see. I see. Yeah, that's

one reason I cannot say that math model exactly fits Amazon, in the sense that there are other sellers providing the same product. And, yeah, I don't have an exact answer. The main reason is that -- so in my model, a consumer can examine only one product, but to take into account the existence of multiple sellers, we have to consider not only the competition, but also we have to alter the assumption of how many products the consumer can examine.

So if multiproducts that I can consider an extension where there are multiple sellers, there are k sellers, they make recommendation, and I can randomly pick two, not just one, and take the whichever better for me, and that kind of competition can actually turn over the -- my main result monopolist, but I think taking into account that there are other sellers in Amazon providing the similar or same product, I don't have a good idea formulating, mainly because I don't know how to think of the consumers' reconciled limited attention with existence of multiple sellers.

MS. JIN: So I have -- I'm wondering how your model changes with dynamics. We know many sellers keep the data in their house for very long time, and if -- I'm reluctant to talk about costume -- Halloween

73

costume this year, but they can use my talk last year or my purchase last year to inform my preference. On the other hand, you could also have consumers learning the quality of the seller recommendation over time.

MR. ICHIHASHI: Yeah. So I think the multiple product with dynamic pricing, so dynamic pricing and the temporal pricing, intratemporal pricing of its consumers is an interesting extension, and I can imagine if the consumer can take into account the impact of today's disclosure in the old future, then I think that -- I believe the similar economic force should arise, namely, consumer can -- are discouraged from providing information if they know information is going to be used in the future.

But I think that another topic I'm interested in is consumer typically know in the very long future, he doesn't know how the information will be used, so once we incorporate something like the time inconsistency or the fact that I don't know in the future how my information will be used, then, you know, there may be something interesting, but that's the case, yeah. I don't know, and I'm interested in it.

AUDIENCE MEMBER: And so my question is also related to Ginger's question. So here you're assuming

74

that consumers control exactly what information is given by the firm, and then in the next round, you're assuming that consumers are not attentive and they can't go and see other products, when in the first round you were sort of assuming they know exactly their value for all the products. How do you counterbalance those assumptions?

MR. ICHIHASHI: Yeah, and that's a very good question. So one is limited attention for product search, and the other is, in some sense, inattention, so a limited attention product search and have attention to control the information. So I think the -- like website like Amazon, putting aside whether it's monopolist, say Amazon, eBay, it's the institutional feature of the websites to show only a subset of the products. I mean, they can never move to a situation where they have alphabetical listing of all the products.

So in that case, no matter how sophisticated about controlling information, it's impossible for me to exhaustively evaluate all the products. So in that sense, I think the -- being able to figure out how to retrieve information doesn't necessarily contradict not being able to find exact product I want by myself, but that point makes sense.

75

Another situation I think -- model may fit is offline transaction, like the car dealers, where I give the car dealer, the salesperson, some preference about my car and get some suggestions and do a test drive. So in that situation, I think it's relatively easy to imagine how to conceal some information. I would not want to dress like this, I would wear cheap clothes, or I may be more reluctant to talk about my preference over fuel efficiency and horsepower if I know that price is very flexible based on the customer. But that's a very fair point.

MR. BRUESTLE: Hi. Steven Bruestle, Federal Maritime Commission.

Sellers often try to induce me to disclose my information by offering me a discount. Have you considered how this would affect your model? Maybe it would be something in between your two cases.

MR. ICHIHASHI: So what you mean is that if you provide information, I will give you a discount.

MR. BRUESTLE: Right, say 15 percent off -- 5 percent off, 10 percent, 15 percent off the final price.

MR. ICHIHASHI: Yeah, right. So if the seller can -- so in the model, if the seller can make the price contingent on what the information disclosed,

76

then as Guy suggested, there is a dystopian situation. I don't know how exactly to incorporate that observation, but the idea that the seller's ability to make transfer contingent on what information disclosure, that's reflected in the market for data part, the very end, in the sense that if the consumer can say no to the transfer and then start to play the original game, he disclose whatever he want, then in equilibrium, the seller makes the transfer, which makes the consumer weakly better off, but consumer is willing to disclose, and the seller can tailor prices.

But in that case the consumer get money, happy to disclose information, and then he's price-discriminated on the product. So he may be -- seller may get a high or low price depending on the informational content, so that doesn't exactly fit what you are saying, but, yeah, I'd have to think of how to get a model explaining that observation. But I agree that that's a pretty common phenomenon.

MR. BRUESTLE: Okay. Thank you.

MR. ICHIHASHI: Yeah.

MS. CARLSON: Any other questions?

(No response.)

Thank you.

MR. ICHIHASHI: Thank you.

77

(Applause.)

MS. CARLSON: So we will take a short break for coffee and conversation. We will reconvene back here at 11:20.

(End of session.)

78

KEYNOTE ADDRESS:
HOW EFFICIENT IS DYNAMIC COMPETITION?
THE CASE OF PRICE AS INVESTMENT

\- - - - -

MR. ROSENBAUM: Hi, everyone. Before I turn it over to Julie to introduce our keynote, just two brief announcements. One is we have copies of the papers that are being presented today on the back table, so if people are interested, please feel free to take those.

And the other announcement is, unfortunately, we don't have WIFI. It's not working in the building. If for some reason you don't have internet and need it for some urgent reason, please come talk to me or Nathan and we will get you hooked up with something.

So I'll turn it over to Julie.

MS. CARLSON: Well, welcome back from the break.

It is my pleasure to introduce our keynote speaker. Dr. David Besanko is IBM Professor of Regulation and Competitive Practices at the Kellogg School of Management at Northwestern University. Dr. Besanko's research covers topics relating to industry dynamics, industrial organization, and the economics of regulation. He has over 50 articles

79

published in leading professional journals in economics and business, including in Econometrica, American Economic Review, Quarterly Journal of Economics, and Review of Economic Studies. Dr. Besanko is a Northwestern University Kellogg graduate, having received his Ph.D. in managerial economics and decision sciences in 1981. Please join me in welcoming Dr. Besanko.

(Applause.)

MR. BESANKO: Thank you, Julie.

I want to thank the Bureau of Economics for asking me to be on the scientific committee. You know, what we did, Ali and Katja and I, was really kind of the tip of the iceberg to all of the work that the economists here in the Bureau did. We each read about a dozen papers and, from those, put together the program, but before we got to that point, the economists here at the Bureau had read, what, over 150 papers, something like that, and so there's a lot of intellectual heft that's behind this conference.

When I joined the faculty at Kellogg 25 years ago, I joined the strategy group, and occasionally people would ask me, what is someone who is doing economics of regulation doing in a strategy group? And I would sometimes say that, well, my goal

80

eventually is to work on research that's at the intersection of competitive strategy and economics and regulation. I'm not quite sure that I've ever got there, but I do hope that this stream of research that I've been doing with Ulrich Doraszelski and Yaroslav Kryukov, and in the past with Mark Satterthwaite as well, is kind of inching in that direction.

So today I want to talk about this topic that we have been working on for the last eight to ten years, which is trying to understand markets where price serves as an investment, and, in particular, for the paper that I'm going to talk about today, we're going to be focused on the welfare economics of such markets.

So what I mean by price as an investment is that there are a lot of interesting settings where it makes sense for companies to sacrifice current profit by setting a low price in order to generate volume that allows it to build some sort of dynamic resource. It's an investment in a dynamic resource that's at the heart of some type of competitive advantage. So an example might be accumulated know-how when there's a learning curve or an installed base of consumers when you have network externalities or switching costs.

And I think these are interesting settings not just from a competitive strategy perspective, but they're also interesting because they give rise to some policy questions. For example, in competition policy, you know, how should we be thinking about pricing below cost, when pricing below cost might really well be at the heart of a strategy to exploit the learning curve, for example? Or in industrial policy, how are we to think about subsidies, which are intended to help an industry take off in the face of these kinds of potential dynamic advantages?

And our view in this paper is that, you know, you really need a well formed understanding of the welfare economics of competition, of competition for the market in particular, in these settings in order to have a useful conversation about policy. So, for example, if unfettered dynamic competition for the market is fairly efficient, then perhaps there might be a relatively big downside to subsidies if you're trying to get the market to take off.

Now, you might say, well, hey, there's really nothing to see here; let's just move along. Maybe the welfare economics of price as an investment is actually fairly clear-cut. Yes, you have this jostle for advantage that results in low prices, at least in the short run. That's good, you might imagine, for

consumers and society, and unlike a rent-seeking model, for example, the value here is not destroyed but presumably transferred to customers through low prices, and so it might be fairly clear-cut that competition for the market, when price serves as an investment, is likely to be pretty good for welfare.

But then when you think about it a little bit, you can see actually that there could be two sides to this. So on the one hand, competition for advantage through price could offset the market power that might typically arise in an oligopoly market. The competition for advantage might actually hasten the investment in these valuable resources, like cumulative know-how. And so you might imagine that the competition is likely to be at least if not inefficient, relatively inefficient.

On the other hand, we need to keep in mind that prices that are too low may actually cause deadweight losses, just as prices that are too high can cause deadweight losses, and in the dynamics of competition for the market, there might be an interplay -- perhaps a dysfunctional interplay -- with various problematic entry and exit dynamics. You might have coordination failures, for example, with respect to entry. You might have wars of attrition when it comes to exiting

the market.

And then, finally, what we have found in other work that we've done is that when you have price as an investment, you can get pricing dynamics that look a lot like traditional notions of predatory pricing, where a firm prices low, a rival exits, and then the firm raises its price, and the long-run market structure turns out to be a monopoly. So it seemed to us at least to be an open question, how efficient is competition for the market when price serves as an investment, and that's really the focus of my talk today.

So the agenda here is to use what we call quantitative theory in the Ericson and Pakes 1995 tradition to assess essentially how efficient competition for the market is when price serves as an investment. So we're going to analyze a discrete time stochastic gain. We are going to compute equilibria over a wide swath of parameter space to highlight the implications of the model for industry dynamics. We are then going to assess the deadweight losses that arise. We are going to assess them against what we hope are interesting benchmarks. And then we're going to anatomize the deadweight loss; that is to say, we are going to decompose it to try to figure out what's

going on.

I actually have two objectives for this talk. The first is to say something, I hope, that's interesting about the welfare economics of competition when price serves as an investment. The second objective is to illustrate what I think is a research question for which quantitative theory is really well suited.

So, you know, we know in dynamic Markovian models, in the spirit of Maskin and Tirole, for example, that pretty much anything can happen. We want to push a little bit beyond that here because we're not just interested in what happens, but we're interested in magnitudes and the patterns that reside in what happens, and we think that this quantitative theory approach is a useful way to do that.

So we focus on one application in the paper, that is, learning-by-doing. This is both economically and empirically important. You can look at Levitt, List, and Syverson, for example, and the dozens and dozens of references in that paper to see the importance of learning-by-doing, and actually to see a very nice discussion of its role in endogenous productivity growth.

We know that learning-by-doing has given rise

85

in the past to interesting pricing and market structure dynamics. You can see this, for example, in Benkard's 2004 paper on wide-body airframes. You can see it more recently in a nice paper by Reichelstein and Sahoo on solar panels.

And learning-by-doing finally I think is important and interesting to look at because the policy implications, to quote Peter Thompson in his recent handbook chapter in The Economics of Innovation, is complicated. I am not going to read the quote, but he talks about how there are complicated issues around both competition policy and industrial policy that arise when you have learning-by-doing.

So let me outline the model for you. I'll do this fairly briefly. So we're going to look at, in this paper, a discrete time/infinite horizon/stochastic game. This is, by the way, the framework that we have used in a variety of papers that we've done over the last eight to ten years.

So in the model we have the action, and the time period is going to be broken down into two phases. There's a price-setting phase and then an entry/exit phase. A state for a firm in this model is the firm's cumulative experience, except for when that

86

state takes on -- that state variable takes on a value of zero, which signifies that the firm is outside the market. So basically the state space in this model is a pair of states, because we have, at most, two firms.

We're going to imagine here -- and this is actually, I think, an important assumption in this model -- that learning is proprietary. So the only way that you gain cumulative know-how is by selling stuff and producing it, and so there's no way that you can catch up to a firm that has more cumulative know-how than you do, other than to sell more stuff. And if you're an entrant outside this market, we assume that you have to start at the top of the learning curve. We can talk more perhaps during Q&A about what happens when that assumption gets relaxed.

We have an entry/exit phase, as I said. If a firm is outside the industry, it gets a draw of a setup cost that's in a certain distribution with a certain expectation. If the firm is an incumbent, it gets a draw of a scrap value with a distribution as well and an expectation. Both those expectations are parameters in the model, as well as the support of the distribution.

And these setup costs and scrap values are privately observed, and so from a rival firm's

87

perspective, it's as if its rival is following a randomized strategy, and so we describe exit and entry behavior through an exit/no entry probability, denoted by fee.

In the pricing phase, here's the Bellman equation for -- this is for firm one, if firm one is in the industry. This gives me a chance to talk about a couple of model primitives. There's a marginal cost which determines -- which depends on the rate of learning, which is captured by the progress ratio, Ro, where higher values of Ro correspond to faster learning. There's a demand function, logit -- the demand function is given by logit demand, the key parameter in the demand function -- actually, the two key parameters in the demand function.

One is Sigma, which is the degree of horizontal differentiation, with zero being no horizontal differentiation, perfect substitutes, and as Sigma gets bigger, these products become much more differentiated and are close to being independent demands. And then there's P0, which is essentially the marginal cost of the outside good. The price of firm n is denoted by Pn as a function of the state, and then U here is the continuation value after the pricing phase.

88

So let me talk a little bit about the equilibrium pricing condition that comes out of the first-order condition. So there really are three pieces to this condition. There's first a piece that reflects static profit. That's sort of the usual marginal cost plus markup. Then there's something that we call the advantage building motive. The advantage building motive is essentially the marginal value -- the marginal future value of improving one's own competitive position. This kind of term would arise in a monopoly model, for example.

And then there's a term that would not arise in a monopoly model, and that's the advantage-denying motive. This is the marginal future value of preventing your rival from improving its competitive position, and this is a term that not only doesn't arise in a monopoly model, it does not arise in the social planner's model that I'll talk about in a moment.

This advantage-denying motive is interesting. Similar terms, analogous terms arise in other papers on learning-by-doing, other papers on network externalities, switching costs, and habit formation. So this is a term that kind of goes beyond this particular application. In this application, you can see that the advantage-denying motive is actually weighted by the diversion ratio, so in an environment with a low diversion ratio, the advantage-denying motive will be less important.

The advantage-denying motive in our 2014 paper, we talked about how the advantage-denying motive is a really important reason why you get what in a moment I'm going to call aggressive equilibria, equilibria that look like predatory pricing.

Let me talk a bit about our computational approach. So we're going to focus on symmetric Markov perfect equilibria, and we're going to compute them. We're also going to compute the first-best planner's problem as well. The planner's problem is to maximize total surplus, taking into account entry costs and scrap values for exit, and we're going to do these computations by -- when we -- as we vary four key parameters: the learning rate Ro, the product differentiation parameter Sigma, the expected scrap value, x-bar, and the marginal cost of the outside good, P0. We are going to use the homotopy method that we talk about in our 2010 paper to do this.

What we essentially do is we look at six two-dimensional slices of parameter space, and I want to say just a little bit about the ranges that we

choose in our computations, because in the paper, we report a lot of frequencies. You know, this percentage of time of all parameterizations, this is what happens. So we had to really be mindful of how we thought about the parameter choices, because we're doing, like, lots and lots of computations here.

So we tried to make the ranges of these parameters, when possible, to reflect their natural economic values. That would be clearest in the case of Ro, which ranges from zero to one. We want to essentially ensure that we have some interesting economic environments, so we chose the range of X-bar to ensure that whatever its value is, that there was always some degree of sunkness with respect to entry, that entry costs were always to some extent sunk.

We tried to span interesting economic environments, so the range for Sigma is going to essentially map us from perfect substitutes to essentially independent demands. And then, finally, what was perhaps most difficult was figuring out what the upper bound should be for those parameters, in particular P0 and Sigma, that had no natural upper bound. We had to put some limits, after all, on the number of computations that we can do, because it's computation -- this is a computationally expensive

model.

So what we tried to do in choosing those upper limits is to avoid representing essentially identical economic environments. So the acid test that we used was, well, if we increase Sigma a little more, does it change things very much? And if it doesn't, then that would be outside that upper bound.

So we ended up doing computations for a little over 2000 distinct parameterizations. That resulted in about 68,000 different symmetric Markov perfect equilibria. Some parameterizations had hundreds of MPE, what my colleague Mark Satterthwaite refers to as the rat's nest of equilibria, and so I'm going to give you -- show you results over the space that we examined.

So the first thing I want to talk about is a typology of equilibria. So the equilibria tended to be one of two types, what we call an accommodative equilibrium and an aggressive equilibrium. These equilibria, for the same parameterization, involved quite different MPE policy functions and implied oftentimes quite different market dynamics and performance.

Let me give you an example for one particular parameterization. So this is a parameterization that

actually gave rise to three symmetric MPE. By the way, we don't know for sure whether we can compute all the MPE. I mean, we do our best to find as many as we can, but we -- I can't assure you, and we don't have a theorem that tells you, that we have found all of them.

This particular parameterization involved three MPE. In the aggressive equilibrium, let me tell you what the modal dynamics in that equilibrium look like. Both firms essentially entered an empty industry almost right away. Then they battled furiously on price, and at some point, one firm gains a cost advantage, and at that point, there's a positive probability that the rival in equilibria will exit. That exit ends up taking about four or five periods. When exit occurs, the remaining firm will raise its price up to a level that equals approximately the monopoly price corresponding to the marginal cost at the bottom of the learning curve. If you were to look at that kind of, well, what would that look like in the real world, it would resemble -- it would sort of resemble kind of traditional notions of predatory pricing.

The other MPE, actually the second of the three MPE, is an accommodative equilibrium -- I should mention that the third MPE, which I'm not going to talk about, is sort of in between these two. The third equilibrium -- the accommodative equilibrium involves, again, both firms entering right away, virtually. One firm temporarily gains an advantage, moves down its learning curve a little. The rival, though, stays in the market. It tries to make sales. Eventually it begins to make sales, and eventually it catches up with what had been a temporary leader of the market.

And then beyond that point, the two firms march their way in tandem down the learning curve, and they do so charging the duopoly price that -- the Nash equilibrium price that roughly corresponds to the marginal cost at the bottom of the learning curve. And you can see that the performance of these equilibria are quite different. In the long run, in the aggressive equilibrium, we virtually have one firm, an expectation. In the accommodative equilibrium, we're almost certain to have two firms, very different expected long-run prices.

By "long-run" here, I mean imagine how the transient distribution implied by the dynamics implied by the equilibrium policies, imagine how that goes in the limit, and then take expectations over that

distribution, and you can see the expected time to maturity; that is to say, to get to the bottom of the learning curve is very different in those two equilibria.

So this distinction between aggressive and accommodative coincides closely, although not perfectly, in those situations where we have multiple equilibria, the equilibria that have the lowest deadweight loss and those that have the highest deadweight loss. So I am going to use the terms "best equilibrium" and "worst equilibrium" to correspond to a case where we have multiple equilibria, and there is a difference in the deadweight losses, which there always is.

So here's what we get. Now, the deadweight loss numbers themselves actually don't mean anything as absolute magnitude, so it's useful to compare them to some benchmark. The benchmark that we use is what we call industry value added. Industry value added is essentially the total surplus that arises in this market from the planner's problem, from first-best surplus maximization, minus the surplus that would arise in an empty industry when the only thing that you have available to consumers is the outside good. So the accommodative equilibrium's relative deadweight

loss is about 4 1/2 percent. The aggressive equilibrium deadweight loss is about 13 percent.

And then here are some benchmarks. For example, in a dynamic model, if we essentially force firms to be myopic, that is to say, we turn off the investment rule of pricing, the deadweight loss becomes about 16.7 percent. If we have a dynamic model where we, in effect, turn off noncooperative behavior -- that is to say, we allow firms to collude on price but still act noncooperatively in terms of entry/exit behavior -- the deadweight loss is about 16.4 percent.

If you turn off both of these, turn off the investment rule of pricing and noncooperative pricing, the deadweight loss is 28 percent. And then with full collusion, collusion on both price and on entry/exit behavior, the deadweight loss is about 14 percent.

So a couple of observations. One is that there's nothing in the primitives here that suggest to us that the deadweight loss should be in any sense low, and the other thing that's noteworthy is that turning off the investment rule of pricing is actually slightly more damaging than turning off noncooperative behavior, which suggests that the investment rule of pricing might be a strong force in this model for

efficiency.

So here's some data on the -- summaries of data on the deadweight loss for all parameterizations. This is -- the first table is relative deadweight loss for all MPE. The median relative deadweight loss is about 7.7 percent. For the best MPE, 5.7 percent. For the worst, 9.2 percent. And in the majority of parameterizations -- and in some cases a long majority of parameterizations -- these deadweight losses are less than 10 percent.

Here's another benchmark we can compare the deadweight loss to some what we think are interesting counterfactuals. So, for example, if we turn off the investment rule of pricing and we force firms to essentially be myopic and we look at the ratio of the deadweight loss in that model to the deadweight loss in an equilibrium, the median of that ratio is 1.78, and the percent of parameterizations where that is bigger than 2 is about 44 percent. The deadweight loss relative to collusion looks a little bit lower, but still the ratio there is well over 1, 1.44.

Here's another view. This is showing you pictures of each of the six slices that we take where we've shaded in higher relative deadweight loss in darker and darker colors, and if you glimpse hard enough and stare at this a while -- well, actually, you don't have to stare at this for very long to see that anything can happen. There are no unambiguous comparative static results in this model with respect to these parameters, at least.

If you stare at this a little bit, you can begin to see that there is a tendency, although it doesn't happen always, for the deadweight loss to be lower as the learning rate gets closer to zero, as learning gets faster. And by the way, these are relative deadweight losses averaged over all types of equilibria.

So some tentative observations. The best equilibria, which are usually accommodative, seem reasonably efficient. The worst equilibria, which are usually aggressive, are not great, but they're still more efficient than if firms ignore the investment rule of pricing and somewhat more efficient than if the firms colluded. And finally, as I mentioned, faster learning, lower progress ratio, does seem to tend toward a lower relative deadweight loss.

So dynamic price competition seems reasonably efficient or at least not too inefficient, even though there are nontrivial distortions that arise in equilibrium. There are too low prices in some states.

There are almost always too many firms in the short run, whatever the type of equilibrium is. There's overentry. There are sometimes, especially in accommodative equilibria, too many terms in the long run, so there's underexit. And the learning is too slow relative to what a social planner would like to achieve.

So why is this going on? What is at the heart of what seems to be a relatively efficient market outcome, or at least reasonably inefficient or not too inefficient market outcome, but yet with these sorts of distortions?

So what we do is we try to anatomize the deadweight loss. So just to remind you here, the deadweight loss is going to be the difference between the expected NPV of total surplus that arises in the planner's problem, which is the maximum level of total surplus, and the level of total surplus that arises in equilibrium.

The deadweight loss, if you think about it in this model, is really shaped by three things. One is that the equilibrium policy function can differ from the first-best policy function with respect to price. The equilibrium policy function with respect to entry/exit can differ from the first-best policy

function with respect to entry/exit. And finally, those two behaviors can imply different market dynamics.

So, in other words, statewise, the deadweight loss is going to be shaped by differences in static surplus. It's going to be shaped by differences in receipts minus outlays from entry/exit behavior. And it's going to be shaped by differences in the likelihood that the industry tends to evolve toward inherently high total surplus states. So we basically take that intuition and we decompose the deadweight loss into three pieces.

There's what we call the pricing distortion, which captures the expected value, in effect, discounted over time, in statewise differences in static surplus. There's the entry/exit distortion, which captures differences over time and expectation between differences in receipts and outlays from entry and exit -- exit and entry. And then, finally, there's the market structure distortion, which captures differences in the way in which the industry evolves over time.

So real quick statistics on the regularities. There's a positive -- typically positive pricing distortion which says -- which is a sign of two

things, actually, that are intertwined to some degree. It tells us that there's a lot of market power going on or there's some market power going on and there's also an inefficiency in which these firms are using price as an investment. There's a positive entry/exit deadweight loss, which tells us that firms in equilibrium tend to have higher outlays for setup costs and lower receipts for scrap values.

And there's, interestingly, a negative deadweight loss component for market structure, which tells us that the equilibrium tends to place more mass on high-surplus states than the planner's solution does, which is telling us, we think, that the gains in this model typically from product variety are outweighing the losses from too slow learning.

So why is the best equilibrium reasonably efficient? There actually are two reasons for this, which we capture in -- the first of which we capture in a proposition, which basically places a bound on the static distortion on a state-by-state basis, and we argue that this bound actually has bite. As incumbent firms move down the learning curve, the part of this bound that involves D0 is going to go down faster than the square of the margin, because essentially what's happening, as the firms move down the learning curve, is they're really marginalizing the viability of the outside good.

As they become more cost-efficient, they're facing less competitive pressure from substitutes, and the industry demand in this case is becoming less price-elastic. In effect, what's happening is that the Harberger triangle is being squeezed. That's the first reason.

The second reason is that for intermediate levels of product differentiation, the accommodative equilibria tend to have more firms -- two firms, in particular -- in a market, whereas the first-best solution tends to have one firm in the market, and that tends, on the downside, to make the deadweight loss component from entry/exit to be positive, but it also serves to reduce the market structure component, and in some cases that reduction can be large enough that, in fact, that market structure component becomes negative and offsets the entry/exit distortion.

And we actually show in the paper that the gross benefit from product variety is going to be enhanced as learning economies strengthen, and that works to limit what we call the nonpricing distortion, which is the sum of the entry/exit distortion and the market structure distortion.

Why are the worst equilibria not too inefficient? Well, one reason is that these equilibria tend to evolve very quickly toward monopoly, and when these aggressive equilibria arise, they tend to be in circumstances where the first-best solution is actually to have one firm in the market.

In addition, we also show in the paper that the monopoly pricing distortion is bounded, and we argue that this bound actually has bite, first of all as firms move down the learning curve, and secondly, it has bite in those circumstances that actually give rise to aggressive -- or an important set of circumstances that give rise to aggressive equilibria; namely, when there is not very much product differentiation in the market.

So, wrapping up, dynamic price competition, we conclude in the paper, is for sure not fully efficient, but it's reasonably so. There's reasonable efficiency despite equilibrium policy functions that differ very much from the first-best policy functions. We conclude that learning-by-doing plays an important indirect role in containing these inefficiencies. In the best equilibrium, it contains the pricing distortion by working to marginalize the outside good. Despite overentry in the best equilibria, the learning

economies essentially enhance the value of having too many firms in the market. And in the worst equilibria, the learning economies help the bound on the monopoly pricing distortion have some degree of bite.

What are the implications for policy? Well, in 34 seconds, it's difficult to talk about all of them. We hope that there are some interesting ones. I'll mention my own view about this, though. I don't see this as a paper that would justify laissez-faire. You certainly would want to have -- this is obvious, I think -- but you certainly would want to prevent collusion in this kind of market. You probably would want to prevent markets -- we want to prevent firms from engaging in exclusionary behavior that would prevent this kind of competition from occurring in the first place.

You may want to think about in this kind of market things that you could do to make learning less proprietary. So, for example, limitations on noncompete clauses that might make it difficult for workers that have knowledge embedded in them from moving firm to firm. I think an interesting direction going forward with this research agenda is to explore in more detail some of these policy implications.

Maybe the one that I'm especially interested in is doing something around industrial policy.

Thank you.

(Applause.)

MS. CARLSON: So we have time for maybe one or two questions, if there are any questions for Dr. Besanko.

MR. BESANKO: Yes, Eric.

MR. RASMUSEN: (Off mic.)

MR. BESANKO: Identical cost functions -- the question was, in the model, are there identical cost functions except for the setup costs. They are -- they are identical de novo, but once firms start to move down the learning curve at different rates, then those marginal costs become different.

MR. RASMUSEN: Oh, okay. You have got a lot going on already.

MR. BESANKO: So, yeah, there's an endogenous degree of asymmetry between these firms.

MR. RASMUSEN: Yeah. But for policy purposes, it would be important to think about where you don't know your marginal cost in advance, because that's a usefulness of the war of attrition afterwards --

MR. BESANKO: Absolutely. So I think -- I don't have time to talk about sort of quantitative theory writ large, that I think would be an interesting discussion, but I do think one interesting question for quantitative theory -- whether our computing abilities are up to the task is less clear -- would be to have models where you have asymmetric information, where other states would include, you know, beliefs about information that the other parties in the game have. I think that's a problem. Asymmetric information in these models, besides kind of the simple asymmetric information that we have around entry costs and scrap values, I think would be a good direction to go.

MR. RASMUSEN: That's too hard for you, but what you can do is symmetric unknown marginal costs, where everybody finds out once you get in.

MR. BESANKO: Yes, absolutely.

MR. BRUESTLE: Steven Bruestle, Federal Maritime Commission.

I'm particularly interested in your policy implementations as to whether or not we should increase or try to increase or decrease learning-by-doing. So it seems like there's forces that could go either way. Do you think, in general, we want to increase or decrease learning-by-doing in firms?

MR. BESANKO: That's -- so I don't want to speculate too much about that. We do find this general tendency that faster learning makes the market more efficient, lower deadweight losses.

MR. BRUESTLE: Okay.

MR. BESANKO: There's something to be said for things that you could do outside the model here that would lower progress ratios generically.

MR. BRUESTLE: Hmm.

MR. BESANKO: You know, so how do workers learn? How do they learn -- how do they learn more quickly? How do firms learn? So I think there's a -- that's all a black box in our model that I think would be interesting to kind of try to break open a little bit.

MR. BRUESTLE: So have you thought of maybe looking at maybe a government endogenously setting learning-by-doing and seeing what level they would want to set?

MR. BESANKO: I would love to have, in a model like this, the Government as an actor. Actually, there is a very -- we have not done that. I would love to go in that direction. There is a very interesting paper that's going to be coming out in the JPE by Mermelstein, Nova, Satterthwaite, and Winston,

which uses this kind of technology, if you will, for merger analysis. They don't look at learning-by-doing, but they actually look at capital accumulation, and in their model, the antitrust enforcer is an active player, and so I think that's a useful direction.

We thought a little bit about that with respect to an enforcer who was going to be policing things that could be considered exclusionary, but I think that's a useful direction, maybe especially for, you know, effecting learning-by-doing.

MR. BRUESTLE: Thank you.

MS. CARLSON: Thank you.

(Applause.)

MR. WILSON: Thanks very much, everyone. If you are interested in lunch, there should be things set out to my left, along the back wall. Thank you very much. We will reconvene in about 30 minutes for the afternoon sessions.

(Whereupon, at 11:59 a.m., a lunch recess was taken.)

AFTERNOON SESSION

(12:31 p.m.)

PAPER SESSION:

THE EFFECT OF PRODUCT MISPERCEPTION ON ECONOMIC OUTCOMES: EVIDENCE FROM THE EXTENDED WARRANTY MARKET

- - - - -

MR. ROSENBAUM: All right, everyone. I hope everyone's enjoyed their lunch. We're now going to get started with the afternoon sessions. So the first thing we have up is a paper session chaired by myself and my colleague Ted Rosenbaum of the FTC.

Our first paper will be by Jose Miguel Abito. He will be discussing the effect of product misperception on economic outcomes.

Jose? Well, maybe we will be taking a slightly longer lunch break for just a minute or so. I hope everyone's having a good day. I appreciate you hanging out inside.

(Pause in the proceedings.)

MR. ABITO: Sorry about that. We were finishing lunch.

So this paper is -- first of all, thanks for accepting the paper, and this work is joint with Yuval Salant from Northwestern and Kellogg, so everybody is kind of -- I graduated from Northwestern, so there is a lot of, like, Northwestern connections here.

Okay, so this paper is about extended warranties. Probably most of you are familiar with extended warranties, but just in case you haven't heard someone selling extended or you haven't encountered anybody trying to sell the extended warranty, so the way we would think about extended warranties is that it's an insurance product that protects you against failure of a durable good. So popular examples of extended warranties are you have extended warranties on vehicles, you have extended warranties on electronic goods. So we are going to specifically focus on TVs in this project, okay?

So what's interesting about extended warranties is actually -- and one that raises concern -- is that typically when you're buying, let's say, a TV, you do a lot of research about a TV with different kinds of brands and, you know, what features they have, but you -- at least, you know, it's rare that you actually think about these extended warranties or even have that as part of your decision-making.

And usually for these products you actually -- even though you're aware of extended warranty, you actually don't know, you know, the terms of extended warranties and specifically the price. So typically a

salesperson, after convincing you that this product is so great and, you know, you should buy it, actually right before you are going to pay for the product, they would say, oh, you know what, it might actually break down and, you know, here's an extended warranty for X dollars, and it's going to cover you for two more -- well, they don't say it's two more years than the manufacturer's warranty. They always say, like, it's three years, okay? So it's usually offered at the point -- all the information that you may have as a consumer actually happens at the point of sale.

So extended warranties are pretty popular. One is that they're very expensive, and if you ask -- like, when I was starting this project, like, whenever I asked my colleagues or talked to them that I work on extended warranties, they would say, well, who the hell is going to buy those extended warranties, okay? So, in fact, the conventional wisdom is that, you know, these are very expensive and mostly useless products, okay? So you even have, like, the Samsungs and -- you know, talking about extended warranties, and kind of like that's the general idea about, you know, the value of these extended warranties.

But despite that, okay, it's very profitable. We weren't actually aware of it. Thanks to Yuval's

student who worked for a consulting company and wanted to get some brownie points from Yuval, we actually found out that it's very, very profitable. And, of course, the companies -- you know, the retailer -- the big box stores didn't want that to be advertised, okay?

So, for example, in the U.S., okay, almost half of Best Buy -- in fact, I think I have seen a number that's more than half of Best Buy's operating income actually comes from extended warranties, and, in fact, the way this -- you know, the way this -- the way these are sold actually may be the reason why these big box stores, if they still exist, are still existing, okay? And the profit margins on extended warranties can range from 50 to 60 percent, okay?

So in the UK, which was relatively more active in terms of investigating the market, okay, they estimated -- when they looked at this market, they estimated that for the top five electronic retailers, they earn roughly, like, 100 million pounds annually, okay? So we wanted to understand this market more, so, you know, we have a lot of preconceived notions of what this market is, but we wanted to go to the data, and, in fact, we were pretty surprised that, in fact, the significant fraction of people actually buy these

products.

For example, one of -- one out of four, like 25 percent of TV buyers actually do purchase extended warranties on TVs. So that's what we saw in the data, and, of course, confirming what we already know, the margins are pretty big, okay? So, for example, on average -- also, so this one out of 24 is for TV, but then you -- it actually -- it's actually the same, so roughly 30 to 40 percent across different product categories, okay? So some products you would think, oh, it might be worth buying extended warranties, but other products, you pretty much think it's of no value.

The margins are big, especially if you compare that to the actual failure rates, okay? So the failure rates was about 7 percent, but then the way -- the price of the extended warranty is roughly, like, 20, 25 percent of the -- the price of the extended warranty is about 20 to 25 percent of the price of the good itself, okay? So that's one question.

Another thing is that -- why we're interested in it is that, you know, it has caught -- because it is very profitable, but at the same time, you know, a little bit dubious in value, competition authorities, okay, or agencies have started or at least caught

113

their attention and actually tried to do something in terms of, like, understanding this market.

So, for example, the FTC, one thing that they have a page talking about, okay, what you should do when you -- you know, when you're faced with a salesperson who's trying to sell you extended warranty, and basically the main message here is try to think first before you actually buy. So they really, like, okay, they -- in the website, they would say, okay, you might actually not benefit from it, okay? Stop, think about that, okay? Maybe it doesn't really need returns or repairs, or, in fact, the potential costs, expected costs are actually pretty low. So, you know, stop first before you think, but that's it.

In the UK, they're more active. They actually did a thorough investigation in the early -- in 2003. What they concluded is that there's insufficient competition, mainly because how these extended warranties were being sold, and there's also -- they did mention there's a lack of information, but mostly they focus on the Competition Act.

And, in fact, around I think 2011, what they did to address that or what they think that could help address that problem is that they forced or they

115

is it something about the buyer himself? Is it that they're very risk-averse and, therefore, they're willing to buy these contracts, okay? Or is it something -- what we are going to explore as what you call probability distortions or, simply said, there's something about how they misperceive or how they kind of distort the decision-making process that they have when they're evaluating the value of these warranties.

So we're going to look at these different explanations and see, you know, which one is more likely explaining it. Then once we have -- and what we are going to see is that these probability distortions is actually driving this business, okay, this market, and -- but it's actually going to be important to understand what actually is probability distortion, what's driving probability distortion in the first place. So we're going to -- we have these two explanations, which is overestimation and overweighting, okay? And we're going to talk about it a little bit more once we reach that question.

And then why do we care about the mechanism, okay? Well, again, it's to actually put scope and the rationale for intervention in the first place, okay? And once we establish that there is some scope and rationale for intervening in this market, then we have

114

required all the big retailers in the UK to actually post their extended warranty price on the website. So there is this price comparison website. When you buy a TV, you will go there, type your TV, and then you see all the extended warranty prices and terms of all the big retailers in the UK. So that's kind of like their way of remedying this apparent problem, okay?

So our paper answers -- tries to answer these three research questions. One is, why is it very -- why is the extended warranty business very profitable? The second is that what drives -- you know, what's the underlying mechanism? Once we kind of understand what's going on, what's the underlying mechanism? And tied to that is that, okay, once we understand the mechanism, can we actually do something about it, okay?

So there's these three questions, first about profitability, and the way we're trying to think about or trying to answer this question is that we're going to explore factors at the buyer and the seller sides in particular. And is it about market power? Is it because of the fact that they essentially have a monopoly on these consumers who are presented with this product at the point of sale?

Is that the reason why it's so profitable, or

116

to think about what tools should we use, okay? We are going to focus on two tools reflecting the fact that we're at the FTC, so we are going to think about competition policies, okay, and also what we call consumer policies, okay, something that addresses more about the decision-making process of buyers, okay?

So let's go to the first question. Why is it profitable? To answer this question, we go to the data, okay? So this is pretty well known data set, at least in the operations/marketing crowd as well, and so it's data coming from a big U.S. electronics retailer. We don't know what it is, okay, but you know how many stores they have, and you can kind of, like, figure out which one it was, okay? So we think it's Best Buy, but, you know, what do we know? So it's a U.S. -- a major U.S. electronics chain, okay?

So we see data from -- oh, so we have data on about 45,000 transactions, okay, and these transactions involve potential purchase of extended warranties. So the data contains everything that's being sold by these retailers, so it's across different product categories, okay? And then we have about 20,000 households, it's a panel, and so the data follows these 20,000 households from 1998 to 2004.

What's interesting about this data is that --

117

so this extended warranty attachment rate is basically how many of those who bought a TV actually bought an extended warranty as well. So across product categories, that's about 29 percent, okay? And, of course, there's variation within categories or across categories. And the ratio between the extended warranty price and the product price is about 24 percent, so they are being priced as 24 percent of the price of the product, okay?

So we are going to focus on TVs because we know a little bit about the failure rates of these TVs, okay? For the statistics of TVs, okay, it's about 27 percent attachment rate, price ratio is 22 percent, and average failure rate is around 7 percent. So the whole paper is going to be focusing on TV purchases, okay?

All right. So although this data is from '98 to 2004, okay, if you actually -- we went back to, like, Best Buy and the other stores, and we looked at their prices. Practically, you know, they're still charging the same high amount, 20 to 24 percent, and actually failure rates have decreased, okay? So it's not that, you know, that a lot has changed, at least in how they're pricing these goods, okay?

So given the data and given our intention in

118

terms of -- or our approach in terms of answering this question, so remember, we want to see whether it's about market power, is it about -- something about consumer decision-making, okay? In particular, for consumer decision-making, we need a model where you have risk -- standard risk aversion, so something that's related to the curvature of your utility, okay, and this notion of distorted probabilities, okay? So we're basically following what's in the literature.

So there's this nice AER paper by Barseghyan, Molinari, O'Donohue, and Teitelbaum. So they look at home and auto loans, and they have this model. So this model tries to explain, oh, is purchases of these insurance contracts driven by standard risk aversion or is it something about the way they're thinking about the probability that you would need or you would use these contracts?

So here let's focus on the utility of not buying extended warranties, so this is where the nonstandardness actually arises. So Phi is the probability of failure, the actual probability of failure. So you're -- the -- typically, okay, so when we're computing the utility of when you don't buy extended warranty or insurance contract, okay, with some probability, Phi, you are going to lose the value

119

of the good or you have to replace it or you have to have it for repair, so you're going to have a -- incur some repair cost, let's say P over there, but with the other -- with the other probability, nothing is going to happen, so you go back and, you know, you have to incur your -- you know, your repair cost.

But if you do buy an extended warranty, so regardless of what happens to the good, you are always covered, but in exchange, you have to pay a price, t, okay? So here we are going to use this model. We are going to estimate this model and basically estimate risk aversion and this function, Omega. So this function, Omega, is what really -- is what the consumer is using when they're evaluating the value of not buying the extended warranties, and so instead of taking the actual probabilities as weight -- as the weight in thinking about their expected utility from not buying, there's actually something going on or something that changes this failure rate and -- or distorts this failure rate, okay, and they're actually evaluating, you know, the relative value of buying versus not buying, okay?

From the seller's side, essentially it's just monopoly pricing of the extended warranty, and this comes from Ellison's add-on pricing model, so I am not

120

going to talk about it that much, given I don't have time.

So the key challenge for identification is being able to -- so identification with respect to the consumers or for the buyers, the key challenge is how can you separate, you know, risk aversion -- standard risk aversion versus probability distortions, okay? So this graph shows, okay, on the Y axis, you have the distortion. Let's say the higher it is, the more distorted it is. So, for example, if the failure rate is 5 percent, the higher it is, you know, the more, you know, they're going to -- they're going to weight that 5 percent by, let's say, 7, then -- et cetera. Then on the X axis, you have risk aversion.

So what these curves show you are iso willingness-to-pay curves, okay? So along the curve, okay, you have the same willingness to pay for a good that has repair cost p and some failure rate, Phi, okay? And each point in this space is just a combi -- is a person, so persons are characterized by a combination of r, the risk aversion, and how much they are distorting the probabilities, okay?

And what this shows is that if you focus on, let's say, the dashed red curve, so that's the willingness to pay for a product with repair costs of,

121

say, pm-prime okay? We don't know -- suppose we know that -- we see the product and we see what the willingness to pay of people or of a person is, okay, but we don't know whether it's a person who has high risk aversion but they're not distorting probabilities that much or is it the person with low risk aversion but they are actually distorting probabilities a lot, okay? So there is this identification problem in terms of figuring out who this person is really is.

So the way we are going to do that in the paper is we are going to look at, okay, another product, okay, that has the same failure rate, but, okay, it has a different loss or has a different repair cost, because if -- so, first, for most utility functions, if they desatisfy this single crossing property where if you change the price, okay, even though these two people have the same willingness to pay for, let's say, product A, if we ask them, okay, how about product B, what's your willingness to pay, we would see that they can't have the same willingness to pay, okay?

Specifically, more risk-averse buyers, okay, will tend to increase, okay, if you give them another product that has a higher loss, okay? They are going to tend to value more the extended warranty relative

122

to the other person, okay, if they're more risk-averse, okay? So in a way the willingness to pay increases faster for the more risk-averse guy relative to the other guy, okay? So that's kind of idea of identification. So we use that in the data, do estimation, and what we find is the following.

This is a bit of a messy graph, and so the red dashed line is 45 -- is the 45-degree line, basically saying if your failure rate is 5 percent, then the way you are going to evaluate that in your brain is also going to be 5 percent, okay? What we estimate -- let's focus on the red curve, okay, and the blue dashed line, which is confidence interval. Basically, one, there's a lot of probability distortion. So, for example, a 5 percent failure rate is going to be essentially equivalent to a 13 percent failure rate, okay? All right, so that's one, okay?

So how do we judge, okay -- before that, yeah -- and so what we find is probability distortions actually drive consumer behavior. So when we estimate these two -- so a model with probability distortions and risk aversion, there's barely any risk aversion, okay? That's what we get, okay? And everything is -- seems to be explained more by probability distortion.

So let's look at the market itself. This is

123

sort of interesting aspect of the exercise. Okay, let's take an experiment where you shut down the distortions, okay, the probability distortions, so you kind of like imagine that there's a way to force people to evaluate the value of the warranty, thinking that the failure rate is the actual failure rate, okay?

When we do that, what's going to happen is, so, we have -- when we're looking at quantities and profits, okay, you see first in the quantities, so this blue dash with circles is -- basically that's (off mic), which is like monopoly and having biased consumers. If you remove the bias, if you remove distortions, okay, it's going to go to this monopoly unbias, so you maintain the market structure, okay, you still can price as monopoly prices, but people are no longer exhibiting the distortions.

What you are going to see is that quantity's going to drop significantly, so about 80 percent, and the consequence with respect to profit is actually very large as well. So just if for some way you can actually influence people's behavior in the sense that they're not distorting probabilities, okay, it's going to drastically change quantities and actually going to lower profits by 90 percent, all right?

124

So it seems that it's really this probability distortion story that is driving the high profits in this market, okay, but it's important to understand what exactly goes on in this probability distortion, and right now it's as if we just, you know, have a reduced-form explanation of why people are doing that, and we saw that it has huge consequences on the market, okay, but, you know, what else can we do, okay?

Well, we need to understand the mechanism, so here we're going to look at two, okay, drivers of probability distortion. One is overestimation, basically people just don't know what the failure rate is, okay? And in this case, giving them information may actually help them, and that might be the way to shut down these distortions.

On the other hand, people may actually overweight failure probabilities in the sense that even if they knew what the failure rate is, they're still not going to decide in that way, okay? They are just going to artificially think, okay, it's a low failure rate, but the way I'm trying to decide it, because the imagery of a failure is so, you know, affecting you, you are actually going to inflate, you know, even though you have that information.

So in that sense it's not clear whether you want to intervene or whether you can even intervene and do something, but from a welfare point of view, you are actually not sure just to respect that type of consumer decision-making or, you know, you want to change it, okay? So it's not clear if overweighting is the mechanism.

However, if it's overestimation, one, there's clear scope on what to do, you give them information, but at the same time, why they -- why they're doing that is because they're -- you know, it's actually a mistake, and, therefore, correcting it is welfare-enhancing both from the consumer and total welfare point of view, okay?

So how do we get -- how do we get at the mechanism? So the other -- the first part was using data from Best Buy or whichever retailer it is, but to actually get the mechanism, you can't rely on just purchase behavior, okay, because you have to somehow, you know, have some intervention in figuring out, okay, what exactly is going on. So what we did is we ran an experiment, okay? So I don't have time to talk about the experiment, you know, but what we find is the following, okay?

Willingness to pay significantly drops, okay,

in the treatments where we give them information. So the experiment basically is that, okay, you're -- you're -- we are told that there's this TV with a certain price, okay? One treatment asks you how much are you willing to pay, and then they -- we also ask what's the likelihood -- what do you think is the likelihood that this TV is going to break down. There's one treatment where we reverse the order. And then there's this -- basically the main treatment, which is to actually tell them before -- you tell them that it's 5 percent and then, you know, you ask their willingness to pay.

So we see that just focusing on the means, but everything is reflected in distributions as well, the one where you give information, the rightmost column, okay, there's a significant drop in the willingness to pay once you say it's 5 percent, okay? All right? So I am going to skip this, okay?

This basically says that, okay, what else is left after you give them information? So we use -- what's the nice thing about this project is that we actually used that identification strategy to design an experiment to precisely get at what we want, okay? So to be able to separately estimate these two things.

So this -- we have a second experiment where we

provide information already, and then we look at -- and we estimate how much are you still distorting probabilities versus is this really risk aversion that explains your willingness to pay. So here we say that, okay, probability distortions are minimal once you give them information, okay?

So given that it's -- I guess I don't have much time, so I will just give a sort of punchline. So, in fact, this is a market where consumer policies are actually potentially more effective. And how is that? Well -- okay, so if you -- if you encourage competition, so suppose you have this price comparison website that everybody -- all of the retailers are going to price at marginal cost, okay, so prices of extended warranties are going to be very low, but if you don't correct the distortion or you don't give them information, then essentially you're encouraging more people to buy this useless product even if, in fact, if they knew better, they're actually not going to buy that product, okay?

So in this case, okay, it might be counterproductive to actually do that, okay? And, in fact, it's more helpful -- more beneficial for consumer welfare to actually address the decision-making problem or mistake rather than

encouraging competition, but, of course, if you have both, then that's the ideal scenario.

Okay. Sorry for -- okay, thank you.

(Applause.)

MR. WILSON: Thanks very much. Our discussant will be Ginger Jin of the University of Maryland.

MS. JIN: Well, thank you so much for having me. It's great to be here.

Okay, let me start by saying that I loved the paper. About ten years ago, I tried to persuade my student to look at extended warranty given its similarly high and abnormal profit; however, I was not successful at all. So this paper really satisfies my intellectual curiosity in a long way by sharpening the question in a policy-relevant context. This also drills down into the mechanisms even after we know we're in the box of consumer misperception.

It also provides a rare case that we can compare competition policy with consumer information policy to see -- kind of run a horse race between the two and see which one will be more effective in addressing the market issues. I really appreciate the creative use of (indiscernible) methodology, both the structural modeling of a very impressive data set as well as the complementary experiment they run to get

129

into the key issues. It also provides a good combination of the empirical facts as well as the theory that's well known in the literature.

So just to summarize the main findings, the first one is the high takeup of extended warranty is mostly driven by consumer misperception. I'm quite convinced by that conclusion. Also, they find that a consumer perception is mostly driven by lack of accurate information and in the failure probability versus some alternative explanations. And the third one is sort of a surprise, but I really feel it's very sensible, where they find that fixing the misinformation is much more effective than fixing monopoly power, and fixing monopoly power alone actually would reduce consumer welfare. This is really speaking to the intersection between antitrust policy and consumer policy that's sort of emphasized the point that we not only should think of them as substitutes, and sometimes they would have these sophisticated interaction effects that actually we cannot think of each one in its isolation.

So I have a few comments and hopefully can help improve the paper. The first one is about product substitution. If I understand the model correctly, the model is sort of thinking, okay, the consumer

130

already decided to buy a certain product. It's just a question of whether you want to buy the extended warranty or not. So in the data, you see individual i buying product j with and without extended warranty, and then you observe another consumer buying probably another product with and without extended warranty.

However, my -- at least my consumer experience is not that I already paid for that TV before I consider whether I'm buying extended warranty or not; rather, I probably have settled down on a model, and then the salesman would tell me the extended warranty, and then I may say, okay, that's a good deal or not a good deal, and then I probably would ask, okay, what's a similar extended warranty price on a substitutable TV.

So in that sense, the model could be sort of -- the alternative model could be that the consumer eyes, looking at multiple products, for each one of them will have extended warranty or not warranty situation. So I wondered what do we miss by ignoring this product-level substitution and only focus on this add-on part?

The policy proposed by UK seems to push the market or at least push the consumers to think about the product and extended warranty as a bundle, right?

131

That's by requiring the firms to post the price not only on the product, but also on the extended warranty at the same time. So if we think of the products as a bundle, then it's sort of different from the structure adopted by this paper. So I think it will be good for the paper to clarify at least what we're missing by not focusing on the product substitution margin.

Relatedly, my second comment is about price endogeneity. So let me see if I understand the identification correctly. They basically assume the perceived probability as a function of the real probability, plus some random variation, okay? And then they look at a pair of products that have the same actual failure rate but different prices, okay? And then they are using the moment condition that the difference between those two products in terms of perceived probability is independent of the price we observe for the product, as well as for the extended warranty.

So this sort of requires the price to be exogenous at both levels; however, I can think of at least a few stories that could violate this assumption. For example, the store may set the price according to their perception of the consumer perception of failure rate. So if, let's say, two TV

132

models have the same actual failure probability, but one is a well known brand and the other is not so well known, maybe new and emerging, then consumers may have different perception on the actual failure rate, and I would imagine that the store may want to price them differently, depending on the consumer reputation about those two brands. So that's story one.

And story two is consumers probably really don't know what's the probability to think about when they buy a TV or a consumer electronics; however, they may use the extended warranty price to try to reverse-engineer the probability, at least I did that when I was a consumer. I'm not sure how successful I was, but if I try to say I look at this extended warranty price, which is 22 percent of the actual product, does that make me think about, oh, maybe the actual probability is close to 22 percent or I compare that with my prior and then decide what to buy? If that's the case, then this price of extended warranty would have the signaling feature that could make this independent assumption violated.

The paper is sort of using, at least in the main specification, using the maximum price of the product as the price, so it's probably not as severe as I'm thinking as the actual price in that transaction for the product; however, I don't know to what extent that sort of alleviate the endogeneity problem.

Okay, I really love the experiments. They have run three experiments. One asks consumers to report their willingness to pay first. The second is -- asks them to report their estimated likelihood of failure first. And the third one is providing the information first. So I would suggest to run a fourth experiment to sort of confirm or probably refute my story that the price might be a signal of extended warranty if you sort of present the price of the extended warranty first and then -- just to see how the subject's going to buy the -- the product or not buy the product, or you can even sort of have a middle question, asking them what's the likelihood given the price they face from the store.

So I have other comments, and they are probably mostly data questions. For example, how do the price vary with each other? I don't know whether the store have kind of constant rate -- constant price ratio between the product price and extended warranty price, or that actually vary across products or over time or across different locations of the stores. And I don't know -- probably given that you don't know the

identity of the store, you probably cannot speak much to whether the store have more sort of salesmen devoted to the categories that would generate more profit in this add-on product.

In the experiment, you have to look at the experiment of likelihood first, that's asking them to predict the failure rate, and then report their willingness to pay, and you sort of interpreted this as a kind of a reminder effect, that you remind the consumers to think about probability, which sort of put them in more disciplined way to talk about their willingness to pay.

I guess probably a related story could be that you forced them to be sort of self-consistent. If I have reported the probability to be 5 percent, it will be very hard for me to justify my willingness to pay to be 20 percent of the actual price. So I wonder whether that could be an alternative explanation.

And lastly, I was fascinated by the fact that the -- what do you call it -- attachment rate, that's the takeup rate of extended warranty, vary a lot by income, and actually income is the only factor that seems important to determine who is buying this extended warranty, and the low-income group have a much higher attachment rate than the low-income group,

so this introduce very interesting question. For example, are low-income households more susceptible to this misperception and whether the firms actually try to take advantage of that differential misperception, for example?

Okay, I guess that's basically my comments. I really loved the paper and hope to see the next version. Thank you.

(Applause.)

MR. WILSON: Thanks very much. I think we have got time for a couple of questions.

AUDIENCE MEMBER: I enjoyed the paper, too. Following up on Ginger's comment, I know with extended warranties it's sometimes argued that the decision-making of somebody who's credit-constrained will be different from somebody who's not. So can you speak to that?

MR. ABITO: In terms of credit constraints, yeah, we didn't include that in the model, but what we can say a little bit -- and this is actually answering also Ginger's discussion -- is we actually estimated -- or it's not in the paper, but we estimated the model for households which are above the median category in terms of income and then for low-income category households. So to the extent that

that somehow related to credit constraints, then -- but I think with credit constraints actually it might -- it might be a little bit more complicated in terms of modeling.

But just to say something about heterogeneity in both risk aversion and probability distortion is we do, in fact, see that low-income households are more distorted in terms of the probability distortion. Yeah, yes.

MR. SWEETING: Andrew Sweeting from the University of Maryland. So in your model you had everyone having a homogenous kind of distortion to their probabilities, and I guess to me this struck me as an area where you might think in terms of a kind of sophisticated naive model, where when you are thinking about kind of the welfare implications, probably the existence of the naives is potentially going to be reducing the retail margin on TVs that charge the sophisticates.

And I actually just wondered, were there different policy implications that might emerge in that kind of setup as compared with the framework that you're using?

MR. ABITO: That's a great question. So one way to also -- so the model we adopt is Ellison's add-on pricing model, which these part one specifications -- one tweak of the model is you have sophisticated and naive consumers. And actually, I would like to answer this question echoing back to one of the things Ginger mentioned, is, okay, what if we think about a model where the consumer's thinking about the TV and the bun -- the TV and the warranty as a bundle.

So in the standard -- in the add-on pricing -- well, in Ellison's model, if that's the case, then extended warranty prices are still going to be set at monopoly prices, but that's going to be competed away. So it's actually not profitable for retailers to do that, okay?

On the other hand, if you have switching costs and unobservability of price, then you are going to have the same -- basically you will have monopoly pricing of the extended warranty, but at the same time, it's not going to affect the pricing of the main good. So in a way that -- it actually can -- or it reduces the incentive of firms to decrease the price of the main good to attract people to buy the extended warranty.

And so in order for that to happen, you actually have to have the right mix of, in this case,

sophisticated and naive. If you have too many sophisticated guys, then maybe, you know -- so it -- the reason why you don't want to reduce the price of the main good to attract people to buy the extended warranty is that, okay, you are going to mostly attract the cheapskates or the naive ones -- or the sophisticated ones, and, you know, they are going to take advantage of the lower price but essentially not buy the warranty. So, yeah, it really -- then maybe competition policy has a bigger role, okay?

MR. SWEETING: Thank you.

MR. BRUESTLE: Hi. Steven Bruestle, Federal Maritime Commission.

Is there more probability distortion for newer products? For example, there could be less word of mouth or experience for newer products.

MR. ABITO: Ah, we didn't check that. That's a great question.

MR. BRUESTLE: It could be a really good natural experiment, too. It could be a good proxy.

MR. ABITO: Yeah. One thing that we kind of did -- we did that kind of addresses that is less about the product but more about the consumer. So we have -- I think I did answer this question. We do actually estimate this with heterogenous preferences,

okay, and one thing about -- one form of heterogeneity that we looked at with consumers is that we have this measure of experience, okay, with a good, so essentially you can say, okay, new products, less experience, and you can kind of map that setting into what we did. So we obviously found that more experienced guys almost not have probability distortions.

MR. BRUESTLE: Okay.

MR. ABITO: So in the sense that, okay, newer goods might have stronger probability distortion. So, yeah, yes.

MR. BRUESTLE: Thank you.

MR. RASMUSEN: Actually, kind of along the same lines, do you know if these products have gotten more reliable over the years in a close enough time, because that would explain misperceptions. I remember when I was a boy, we actually had a TV repairman come to our house to change tubes, and disk drives used to fail a lot, hard drives, and I don't hear about that nowadays.

MR. ABITO: Yeah. So definitely failure rates, for example, for TVs have gone down now. It's about 5 percent. I am not sure I got your second question, given --

MR. RASMUSEN: Well, so that's since 1970, maybe. Since 2000, have they gone down?

MR. ABITO: Well, even just comparing 2004 and now, the -- the -- so most of the repairs here is something with the screen, and the technology for developing more reliable screens actually has improved, so definitely failure rates -- so that's actually what's funny -- or not funny, but the failure rates have gone down, but then the prices are still that high. Yeah, so -- yeah. Of the warranty I mean.

AUDIENCE MEMBER: (Off mic.)

MR. ABITO: Have less distortion, yeah, yeah. I mean, they know how to handle the products, and they kind of see that, you know, these are not -- not --

AUDIENCE MEMBER: (Off mic.) More experienced guys would be more distorted if there's higher demand -- higher probability (off mic).

MR. ABITO: All right.

MR. ROSENBAUM: Okay, great. Thank you.

MR. ABITO: All right, thanks.

(Applause.)

MR. ROSENBAUM: Thank you very much.

(End of session.)

PAPER SESSION:
CONSUMER PROTECTION IN AN ONLINE WORLD:
WHEN DOES OCCUPATIONAL LICENSING MATTER?
-- - - -

MR. ROSENBAUM: Our next speaker is Andrey Fradkin from Boston University presenting Consumer Protection in an Online World: When Does Occupational Licensing Matter?

MR. FRADKIN: All right. So I'm really excited to present on this topic, especially here. So I think occupational licensing laws are an interesting topic in and of themselves, but, of course, they're very much predicated on the understanding what types of information does a consumer have as well, and the internet is changing that to a great extent.

So before I get started, I would like to mention that this is joint work with Chiara Farronato, Brad Larsen, and Erik Brynjolfsson.

So in case you don't know, occupational licensing laws require individuals to obtain permission from the Government in order to perform a particular service, and about 30 percent of the U.S. labor force is affected by occupational licensing laws of some form. And you might naturally think of occupational licensing of doctors, where they are

really high-skill professions, but you also have occupational licensing of medium-skill professions, like electricians, plumbers, or painters, and then even low-skill professions or at least low-risk professions, such as hair braiders.

So to give you an example of one profession that we're going to study, interior painting, so the map of the U.S. shows the states in yellow, those are the ones where the states have specific occupational licensing regulation regarding painting. Furthermore, the states in white, they don't have a statewide regulation, but individual cities might have a regulation. So in Texas, San Antonio would have a regulation about interior painting.

And then across the states that do regulate occupational licensing, there's a lot of variation. So for example in Nevada, you have over a thousand dollars' worth of fees. You have to have four years of experience, which is oftentimes apprenticeship. You have two exams. And by the way, the pass rates on these exams are actually not very high. We were able to check that in a few cases. And then one caveat I should point out is that there are oftentimes cutoffs in these occupational licensing laws. So if the job exceeds, let's say, a thousand dollars, then you need

the license, but otherwise, you don't.

So to go over some basic theory about occupational licensing, we think that licensing has the advantage that it's going to protect consumers, especially if the transaction is risky. On the downside, it's going to be a barrier to entry. So it's going to reduce competition, it's going to raise prices, and increase rents for the licensed pros.

So an important contribution in the theory literature on occupational licensing is Shapiro in 1986, which essentially says that occupational licensing is going to be Pareto inferior if sellers' investment in quality is going to be observed or when reputation accumulates really quickly, which brings us to the setting of an online platform.

So online platforms are going to intermediate a lot of these transactions now, and, importantly, they are going to display online reviews and other information regarding the pro. And this changes kind of how we think about occupational licensing in two ways. First of all, it's going to provide for a new way to evaluate whether occupational licensing laws are doing a good job or not, because previously we weren't able to get transaction-level data in these occupations. And then secondly, reviews might

actually serve as a substitute for licenses.

So in this paper we're going to ask two research questions. Do customers care about occupational licensing status on this online platform? And we find that, if anything, they dislike occupational licenses, and reviews matter a lot more for the consumer choices. And then secondly, we study a more aggregate outcome, which are equilibrium outcomes in terms of the match rates, the prices, and the ratings, as they vary by the stringency of a licensing regime across states and occupations, and we find that more stringent licensing regimes lead to less competition and higher prices on average, but no detectable effects on customer satisfaction, so -- and we don't find strong evidence of benefits of occupational licensing.

Okay, so the rest of the talk is going to be in three parts. First of all, I'll describe the setting and some descriptive statistics. Secondly, I'll describe the individual choices. And thirdly, the aggregate outcomes.

Okay. So the setting is an online platform for home improvement services. It has national reach. It has millions of transactions that happen on it. And I'd like to say that home improvement isn't just, like, an interesting test case for occupational licensing. It's an important profession. Broadly construed, there are over 5.3 million workers in the construction industry, and they -- these are the types of jobs that are unlikely to go away any time soon.

So the way the platform works is that a customer will have a local service need, so maybe I'm looking for a painter in D.C. I'm going to Google "painters near me" or "painters in D.C.," and this platform will be one of the top search results. The customer is going to enter the platform, and they are going to be asked to submit a detailed job request. That might say, how big is your place? Where is it located? What type of paint would you like to use? How quickly would you like this done? And other things you might think about.

Once the customer submits this, pros are going to pay to submit a bid on a particular request for a job. So that's the business model of the platform. The pros are paying to get the lead. And there is a maximum amount of pros that can submit bids for a given customer. And then after that, the customer can choose to hire one of the pros.

So here's a stylized version of the profiles that each pro might have. So we have Interiors by

Chiara Farronato. She has one review. In contrast, Fradkin International Design have ten reviews, but there are three stars average rating, and I'm also licensed to be an interior designer. So that's the important part, is that when the platform has verified your license, that license is displayed in your profile information.

So how does the platform verify these licenses? So, first of all, the pro must submit the license to the platform, and then once the platform receives the license, they're going to take some amount of time to verify it, and the way that they would do so is they would go to the appropriate state-level website, so let's say the Licensing Board of California, and they would go look for that ID number and make sure that it matches up with the pro. And a key for us is that the amount of time it takes a platform to verify the license is quasi-random.

So what are the types of jobs that are available on the platform? Lots of contractors, so general contractors, HVAC contractors, painting contractors, and so on and so forth; plumbers, electricians, home inspectors, pest control and pesticide applicators. So you should be thinking about these types of jobs.

In terms of -- oh, I guess old slides. In terms of the summary statistics, we see that at least in our sample of quotes about 12 percent of the quotes are by a pro that has a license validated at that time, and 14 percent by a pro that has submitted a license to the platform at that time. The typical quote has four reviews and a 4.9 pro rating. So as with other online platforms, the ratings are typically skewed towards five stars. It's the -- it's one out of five stars -- sorry, or it's out of five stars that the rating is.

And then conditional on hire, we see that hires tend to have more reviews and lower prices relative to the quotes, which is, I guess, not very surprising.

Importantly, since we're going to be studying both reviews and licenses, we want to see, do licenses predict the quality of the transaction as measured by the rating that a pro receives from a customer? And so what we see is that just the raw correlation, this is a small positive relationship between whether you have the license validated at the time that you did the job and the rating that the customer gives you once you did that job.

But once we control for whether you've submitted the license, it seems that that's what's

soaking up most of that correlation, and this holds when we add more controls.

Another thing we can do is, even before you submitted the license, you were probably already a licensed pro; you just hadn't gotten around to submitting the license to the platform. So we control for whether you've ever been licensed on the platform, and that soaks up a little bit more of the variation.

And then, finally, the last column is going to have pro fixed effects, and we don't see any change in the types of ratings that you get as you get validated on the platform. So I would view this as generally -- there isn't strong evidence that licenses on the platform are predicting five-star ratings here.

So one thing I mentioned previously was pros might not need a license for certain types of jobs. So here we made two plots for California, one for general contractors and one for painters, where California's law is that if it's over $500, you need a license, and so you can see that both pros that have gotten licensed on the platform and pros that haven't gotten licensed on the platform oftentimes bid more than the $500 limit. So this could mean that they are really licensed, they just haven't told the platform, or it could mean that maybe they're not paying attention to licensing laws. We don't know. Okay. So that's a description of the setting.

Now getting to the study of the individual choices, so the basic type of regression we would like to estimate is the outcome variable is whether the customer hired a particular pro as a function of the pro characteristics and the bid characteristics, and the variables that we're interested in are licensed, price, number of reviews, and the average rating that you get. And just to kind of give you a sense, like, every pro is going to bid a particular dollar -- fixed dollar amount, and it might have some text associated with the bid as well.

In terms of the identification strategy, where you need a separate identification strategy for all these variables. So for the licensing variable, we're going to use the fact that there is this quasi-random amount of time that it takes for the platform to verify a submitted license as being verified and to display it on the site. So we're just going to have a control for whether the pro has submitted a license at the time and whether the pro's license has been verified by the platform.

Secondly, our instrument for price is going to be the distance between the pro and the customer. So

the customer presumably doesn't care where the pro is located, but it takes more time for the pro to get to the customer, and that should be a cost shifter there.

And then, lastly, for the reviews and the average rating, we're going to use the characteristics of the prior reviewers of that particular pro. So if the prior reviewers of that pro tended to be harsh, that they reviewed other pros with lower ratings, that should shift around the ratings of a given pro. And similarly, for the propensity of the customers of the prior professional -- of the prior -- for the customers of the pro before this given quote, if they were more likely to submit reviews, then that should increase the number of reviews of the pro.

Okay. So before getting to that full specification, we're going to do an event study analysis. So here we just include professional fixed effects and request fixed effects, and the licenses that we are going to use for this are going to be occupation-specific licenses, so we are going to throw out business licenses and inappropriate licenses. So like some of these might have, like, an accountant license, but they're not doing the accounting job, so we are not going to include that.

So here are the results of this regression,

where the timing is relative to the time when the license was validated. So we don't see a statistically significant effect at the time of validation in the hire rate, and the variation in these coefficients is very small. So there doesn't seem to be much evidence that customers are paying attention to the validation of the license.

You might say, well, maybe the pros are changing their behavior in response to getting the license validated, and we don't see evidence of that either. This is the same regression where the outcome variable is price.

Okay. So now getting to that full specification where we have license, ratings, and reviews, I'll go through each of the specifications in order. So these are the results from the OLS, and the column I've highlighted includes professional fixed effects and request fixed effects. So we don't see an effect of licensing in the specification. We see a positive effect of average ratings, a negative effect of the number of reviews and of the price. So we need some instruments here.

So the next column is going to add an instrument for price, and we see that the coefficient on price becomes much more negative, exactly as we

would expect. And then the last column is going to add our reviews for ratings, and in that specification what we see is a small negative effect of license verified on consumer choices and positive effects of having higher ratings and having more reviews and a negative effect of price. So in terms of relative magnitudes, the license verified doesn't seem very important to these other variables.

We can also look at the same type of regression where the outcome is going to be whether the customer viewed a quote. So the customer gets a list of quotes, and they don't have to view all of them. So we kind of see more views than hires. And when looking at this outcome variable, we see very similar results in the sense that people are going to be more likely to view a quote from a pro if that pro has more ratings, more reviews, less likely if it's a higher price, and having a license verified actually decreases the view rate.

One thing that we thought might be worth looking at is interactions of license verified with review-rated variables, but we don't see any consistent patterns here. So I am not going to discuss this any further.

Okay. So what does this mean? Consumers might

153

not pay much attention to the licensing for various reasons. One might be they just don't really know what a license is doing, so they don't care. They might not pay attention because they rationally know that licenses don't actually affect quality in this market, or maybe they just assume that everyone is licensed and that's why they don't pay attention.

We don't really have much to say about which of these stories is true, although manual inspection of pros suggested that it was very hard to find licenses for some of them, and part of that is just that the name under which a pro might have registered their license at could have been different from the one that we observed on the platform, but it could be that they're actually not licensed.

All right. So now moving on to the aggregate outcomes, so what we want to know is how licensing stringency is going to affect outcomes on this platform in terms of competition, prices, and quality, and the identification we're going to use is going to be across zip code, across licensing stringency. So think about painters and electricians in California versus painters and electricians in Nevada. If Nevada happens to have more stringent licensing for painters, then we should -- then our regression is going to pick

154

up the effect of that relatively more stringent licensing on painters in Nevada compared to California.

So the regression that we're going to estimate is the following, where the outcome variables are going to be the number of quotes that a given task receives, the quoted prices, whether there was a match or not, what the price of the winning quote was, and then outcome variables such as the star rating and whether the customer comes back to the platform, which are measures of quality. And the observations are going to be at the request, zip code category, and month/year level, where importantly we're interested in the coefficient on licensing stringency.

So that actually raises the next question. How does one measure licensing stringency? So we start with a database that the Institute for Justice has compiled called Licensed to Work, which includes, for a wide variety of professions, the fees that you need to get licensed and exams, the minimum grade and age, education and experience.

We've also -- and this is not in progress, this is in the data -- we've also compiled our own information about general contractors, electricians, and plumbers to augment that data. And then once we

155

observe these variables, we conduct a PCA analysis to create a one-dimensional score of licensing stringency, and we're going to exclude in these regressions states that don't have a statewide occupational licensing regulation for a given profession.

So what are the factors that are correlated with this dimension we've identified? They're going to be fees, exams, minimum grade, minimum age, education and credits but not in years, and then experience in terms of years. So generally most of the factors are positively loaded. Most of the variables are positively loaded in this factor.

Okay. So here's the standard regression that we have. So we see that the number of quotes is negatively associated with licensing stringency. The prices are positively associated with licensing stringency. Then there's no effect on star ratings or whether the customer comes back to the platform.

Now, some of these estimates are a bit noisy, and you might also be saying, well, there's a -- kind of one thing that can be very different between different zip codes and different states. Maybe the types of painting jobs that one does in Nevada might be different from Massachusetts or North Carolina. So

156

what we do next is we are able to control for these very detailed requests, characteristics.

Do you have like a 2000-square foot house? Do you need this type of paint or that type of paint? For each separate profession, using the double ML technique of Chernozhukov, et al. So the basic idea behind that is you split your sample, and for one-half of your sample, you estimate a machine-learning model that predicts both the outcome and the treatment, which in our case is stringency, as a function of all these detailed task-level characteristics, where we use a lasso estimator.

And then on the other sample, you estimate the model that you're interested in, which is the outcome variable on the residualized licensing stringency, and we get the following results, which are now precise but actually very similar in magnitude to the ones that we saw on the previous slide. So having more stringent licenses associated with fewer quotes, higher prices, lower match probability, and no difference in terms of customer satisfaction.

We also tried to do some heterogeneity analysis, where each of these is a profession, and we're kind of -- and most of the heterogeneity points in the same direction as the regression that pools all

157

the categories together, and we're -- and we aren't able to, once again, detect any effect, even at a profession level, on star ratings or whether the customer returns to the platforms, which are quality measures of the professional.

Okay. So what have we learned from this? So kind of a narrow interpretation of this is that we've learned the effect of licensing for digitally initiated transactions, which might be a very small subset of all transactions. Kind of a broad interpretation is that we've learned something about the effect of licensing both online and offline, and that would be true if the service providers online are not systematically different from the service providers offline, and if consumers spend the same effort verifying licenses online and offline, which may or may not be reasonable depending on the particular job that we're thinking about.

So in conclusion, how do licenses and reviews affect customer choices? We find that reviews matter much more than licenses. And how do equilibrium outcomes vary with licensing stringency? We see that more stringent licenses are associated with less competition and higher prices, but no detectable effect on satisfaction. And lastly, it's still very

158

much a work in progress, so I'm looking forward to your comments.

(Applause.)

MR. ROSENBAUM: And Judy Chevalier from Yale will be the discussant.

MS. CHEVALIER: I had a little miscommunication about my slides, so I am going to use Andrey's, and his are better than mine were going to be anyway, so we're good.

Okay, great. Thanks. So let me thank the organizers for inviting me to discuss this paper. I've done work in the past on both occupational licensing and review platforms, and so I'm somewhat jealous that I didn't think, you know, to put those two things together.

And, you know, when I think about this paper, what I think the contribution is, you know, when we are -- we're living in a world with more review platforms than we have ever had before, which means that consumers are observing more things about providers in all kinds of spheres than they were before, and even if we thought the occupational licensing regime was optimal at some point in time -- which, by the way, I'm pretty sure we don't -- but even if we thought that, you know, with more consumer

159

information, presumably the optimal occupational licensing regime would change, and so this paper explores that in I think a very interesting and persuasive way.

Okay. So when I look at occupational licensing, the authors emphasize this motivation, and, you know, we're at the FTC Bureau of Economic Analysis where the charge is competition and consumer protection, and the authors focus on competition and consumer protection, and many of us are IO economists, so that is the place to focus, but I do want to back out a little bit from the motivation here that the authors provide and just remind ourselves what other areas of economics tell us about occupational licensing, right?

So what are some concerns about occupational licensing? Well, one thing about occupational licensing is that it impedes economic adjustment because it impedes worker mobility, right? So there are jobs -- if job opportunities are declining in Michigan and rising in Wisconsin, we would like workers to move from Michigan to Wisconsin, and occupational licensing imposes a barrier on that kind of mobility.

Second, occupational licensing imposes a

160

barrier on economic mobility, right? So we have -- the classic example would be, you know, we have a new immigrant who's very skilled at hair-braiding, and, you know, would be a great hair-braider if only she could meet the licensing requirements, all right? And so the ability to that worker to match to the best job that matches her skills is actually impeded by occupational licensing.

And I remind us of this because usually when we do these kind of welfare things in IO, we're kind of wanting to ask ourselves, you know, do these pluses outweigh these minuses, right? Does the consumer protection outweigh the competition and pricing issues? But here I've just told you two labor market things that are just unambiguously bad, right? So the consumer protection part better kind of hit the ball out of the park if we're going to tolerate a lot of really strict occupational licensing, right? So just a framing to think about occupational licensing in this context.

Okay, so the authors have two research questions, and I want to turn a little bit to some critiques of each, and, you know, let me say I'm going to do the usual discussant thing of nit-picking a little bit, though on net, I am very skeptical of many of these state occupational licenses, and I don't want that message to get lost in my nit-picking here, because I think in general I pretty much believe the results that the authors have come up with.

Okay. So the first result is that consumers on this site are not much impacted in their choice of pros and in the reviews that they ultimately leave by whether or not the pros have verified licenses on the site. Now, one thing I would like to maybe spend a little more time on than Andrey did is this issue of consumer beliefs when there isn't a license posted on the site, and, you know, Andrey said forthrightly that it's not clear what the consumer believes in that circumstance, but what I want to point out is what the consumer believes is probably heterogenous across these various occupation types.

So, for example, when I think about an electrician or a plumber, you know, if I see something advertised as Fradkin Plumbing Company or Fradkin Electrical Company, I have a strong prior that electricians and plumbers are regulated, and they're regulated pretty much everywhere. In contrast, painters, as we saw in the picture, are regulated in some places, not regulated in some places, regulated in some circumstances, not regulated in some

circumstances. My guess is the typical consumer, when going to the site, has a strong prior that the electricians and the plumbers are regulated and have occupational licenses, and maybe has a much more diffuse prior about the interior decorators and the painters.

Now, why does this matter? Well, the situations in which I suspect -- but we can't show -- that the prior and the posterior are similar -- that is, that the consumer has some, you know, understanding that, say, everybody's regulated -- are actually precisely the same situations in which the consumer would care about the regulation status, which is to say regulation of painters might be dumb, and we see that the consumers don't -- you know, the consumers didn't know whether the -- it's possible the consumer didn't know whether or not the painter had a license, and the consumer doesn't care, and they're probably right not to care.

But we can't take from the results that the consumer doesn't necessarily care about the electrician or the plumber, because the consumer hasn't actually possibly been updated that much about the probability that the electrician and plumber is licensed by seeing the license verified on the site,

given this consumer probably thinks that anyone doing business as an electrician or plumber has occupational licensing.

So I would want to just be careful about that idea, that in both of the situations I described, we could find that reviews matter much more than licenses in motivating consumers to hire, but that's not exactly the same thing as saying a consumer would willingly hire an unlicensed plumber, all right? So I think we just have to be careful about the interpretation there. And it might be that there's some things that could be done to try to look at that heterogeneity.

Now, let me turn to the second result, which are the results about the stringency of licensing regimes, and here the authors find somewhat compelling evidence that more stringent licensing regimes lead to less competition and higher prices -- that should be kind of satisfying to us as economists, because we expected that to be true -- and also maybe less satisfying, no detectable effect on consumer satisfaction. So consumers are paying more in the circumstance in which they're in a location with more stringent licensing for the particular profession that they're looking at but also, you know, don't end up

any happier about the job that's been done.

Here I do think we have to pause a little bit to think about the difference between what we think, say, the education in the occupational license is teaching and what are the things that consumers care about or measure. I think the results here are most compelling in situations in which, you know, the kind of things that the license would measure are things that a consumer would sort of immediately be able to detect, right? Then we would in some circumstances expect to see some relationship between satisfaction and ratings.

But there are two kinds of things that the licensing might be teaching, let's say the educational requirements. One might be the kind of things that are observed only down the road, right? So the plumber who comes is nice, he seems to do a good job, I give him a high rating, but does it leak later, right?

And then I think another thing we should think about, which is a compelling case for regulation but, again, would get -- and would give us exactly these results, but wouldn't necessarily make us think the regulations are bad, are any situations in which the education or the regulations are on things that

165

consumers systematically underscreen for relative to social welfare.

So what would be some examples of that? I'm pretty sure some of the training for the pest companies is about safe disposal of pesticide. The consumer might not care about that. It might not affect consumer satisfaction, but if you think poor disposal of pesticide is an externality, we want higher prices, and the consumers won't be happier. Similarly, occupational safety kind of stuff for a roofer would be the sort of thing we might regulate, it might be part of the occupational licensing, but we might expect these effects.

Here again, I think one thing the paper could do, which I think would help a lot, is just a little more color on what the licenses -- you know, what does the educational program look like? And maybe a little more digging into the heterogeneity across the various kinds of occupations. But my guess is that I believe you, that there's pretty compelling results here, that there's some sets of occupational licenses in these building professions or home services professions that probably don't create a lot of value for consumers.

Thanks.

(Applause.)

166

MR. ROSENBAUM: All right, we have time for some questions.

AUDIENCE MEMBER: Hi. First of all, as an economist, I thank you from a very -- professional or occupational licenses are pretty important, you know, I don't think too many universities allow people without Ph.D.s to teach. I wonder what our licenses are good for, but...

So the one table actually that caught my attention is where you ran these regressions, licensing but also reputation variables, and you had these very nice instruments for both price and reputation. Typically what we think about in reputation in online settings is, you know, the reputation score is correlated with other stuff in the listing, and it's, you know, maybe upward biased to, you know, capturing, but when you do the IVs, those coefficients jumped up by two or three orders of magnitude, from very low to extremely, extremely high.

So do you have an explanation for why -- I know this paper is very preliminary, but I've never seen coefficients change that way.

MR. FRADKIN: Ah, I don't -- I've never investigated what causes the jump in terms of examining, like, the details of that, but --

167

AUDIENCE MEMBER: It goes from one basis point to, like --

MR. FRADKIN: My sense is that maybe as you accumulate reputation, you start winning different types of jobs, which -- I don't know, which -- which might be harder to fulfill or something like that. I -- but I -- I don't really -- I don't have a good intuition for that. I'll have to think about it more.

MS. JIN: Yeah, I'm interested in the interaction of the two results you show. If we go with the consumer side result, that consumer did not care so much about licensing, and if -- even if you're not licensed, you can still bid and get selected by consumers on this platform, which sort of is not exactly consistent with the second result, that licensing actually reduce competition, and if I can get in the market, that sort of suggests that whatever the regulator is not enforcing that licensing requirement very stringently; however, somehow it still have this lessened competition effect. So do you have any comment on that?

MR. FRADKIN: Yeah. So I think for a lot of these professions, there are these kind of, you know, $500 thresholds, so it's perfectly reasonable for there to be both licensed and unlicensed pros to be

168

operating and competing with each other, but the fact that some of the -- in order to compete for some of the jobs, you need a license, that's going to affect the entire market structure. So I think that's one reason that you could reconcile these results.

AUDIENCE MEMBER: So here in your results you were showing that there is a correlation between five-star ratings and licenses, but that goes away when you control for reputation and other characteristics. So what I'm wondering is, is it common because people do not update or include their licenses early enough, that -- like if I -- when I'm very new to the system, maybe at that time, when I have my licenses in the system, it will help me, but the value of the license will go down when I have built up my reputation? So if I'm there for two years and I had probably license all the time, I just didn't upload it on the system, it -- after two years, it would not have any effect, but if I had done it right away, it would have had some impact for my future?

MR. FRADKIN: Yeah. So we haven't done that heterogeneity analysis, and we can check for that, but at least when we think about how consumers react to the verified licenses when they're making their hiring decisions, we weren't able to find kind of very

169

smoking gun interaction effects in the amount of experience that you have and whether you got the verified license or not, in terms of how consumers are hiring them, the pros.

MR. RASMUSEN: Okay. Well, that was my question, but I'll ask another one. So you did the interactions and no effect came up between experience and license?

MR. FRADKIN: I mean, it's inconsistent. It's hard to -- like, depending on the specification, we'd get -- we'd kind of get all sorts of weak effects.

MR. RASMUSEN: A lot of ways to do it, yeah. Anyway, I'll suggest something more along the line of Judy's, in that one good reason for licenses would be so you can take them away if the person misbehaves. It doesn't apply so much to painters. They might burglarize your house or something, but otherwise you could have a system where you pay a hundred dollars to become a doctor, and then if you kill somebody, they take away your license and put you on an online list, and I know Indiana -- and there may be other states -- have online all the horrible stories of people you might hire for something.

MR. FRADKIN: Yeah. So I've seen those lists as well. We haven't thought about incorporating them

170

into this analysis, but I'll think about it more, yeah.

AUDIENCE MEMBER: So I'm thinking about this, coming back to the platform outcome variable on the left-hand side, and I think it might be working against you on some level, because if I have a house with many electrical problems, I go to the platform, I find a good electrician, the next time I have an electrical problem, if I'm happy, I'm not coming back; I'm just calling that person, right?

So on some level, not coming back to the platform in the same category is a good outcome, not a bad outcome. So maybe there's a way of breaking apart coming back to the same category versus coming back for something unrelated and seeing if that helps you out.

MR. FRADKIN: Yeah, that's a great suggestion. We haven't done that yet, but that would work.

AUDIENCE MEMBER: This is sort of orthogonal to your research question, but is there any sense in which there are other effects of occupational licensing that -- on the firm side? Like can firms be insured if they are not licensed in certain states, and if you slip and fall in somebody's house while painting, like, most of the time you wouldn't if there

171

was a five-star review, but do you know whether -- how things like that work with licensing?

MR. FRADKIN: So my understanding is that some but not all licenses require insurance. I think it depends on the profession. Oftentimes, the pro will signal in their profile text that they are insured, but it doesn't happen, like, an overwhelming amount of the time. So I -- that's -- that's something that we are going to have a hard time studying in this paper, but I agree that that could also be important, and especially to the extent that getting insurance might be more or less difficult for certain types of individuals.

AUDIENCE MEMBER: So it's not -- okay, yeah, we can talk about it.

MR. FRADKIN: Yes.

MR. ROSENBAUM: Thank you.

(Applause.)

(End of session.)

172

PAPER SESSION:

DIAGNOSING PRICE DISPERSION

MR. PETEK: All right. We are ready to get started again. Our last paper session of the day is chaired by Ali Hortaçsu. I'm Nathan Petek. I'll be introducing the speakers. The first speaker is Matt Grennan, who is going to present Diagnosing Price Dispersion.

MR. GRENNAN: All right. I want to say thank you to the organizers for putting on this great conference and the organizers and Ali for including us on the schedule today.

So this paper is part of an agenda that Ashley Swanson -- sitting over here -- and I have on thinking about the prices that hospitals pay for medical technology inputs, and it turns out that if you look at the exact same input at the exact same point in time sold by the same vendor and then look across hospitals, they will be paying quite different prices for that same thing.

This figure right here is representing that by looking at a given product, looking across hospitals at a point in time, and taking the coefficient of variation -- so the standard deviation over the mean -- for that particular product, and then this is

173

summarizing that over products for a bunch of different product categories. So each of those dots or squares or triangles you see there is the mean of that measure for a product category, and then the bar you see around that is the inner quartile range among products in that category.

So the take-away from this is that there's a lot of this price dispersion across hospitals or across buyers for the exact same thing, no matter if you're looking at some of these PPIs, or physician preference items, you know, these are like the high-tech medical items, you know, stents, pacemakers, hip and knee implants, ranging, you know, from those on to kind of more commoditized items, you know, like needles, surgical gloves, and so on.

And these price differences are pretty meaningful to a hospital's bottom line. So hospitals run on pretty thin operating margins, so the average AHA survey margin in 2013 was about 3 percent, and these -- and the supplies that are in the database that we're going to analyze today represent about 23 percent of hospital operating costs. So if you do the back-of-the-envelope math here, you're moving one standard deviation, and all of these supplies would be kind of equivalent to going from the average to going

174

to, you know, redlining it.

So these are pretty meaningful differences, and so we want to look at what's underlying some of this variation across hospitals, and then we want to look across these very different product categories and see the extent to which those underlying features may be similar or different.

And so why does this kind of law of one price tend to fail here? So there could be many reasons, right, many of which a lot of people in the audience here have studied, and, you know, one would just be there is some sort of brand preferences, right? So these are differentiated products, and maybe the preferences over these differentiated products are different among physicians or providers at different hospitals, and that's some of what we're seeing here, right?

Another would be a variety of explanations on the supply side, so maybe distribution costs may vary somehow, and some of that's what we're seeing; maybe, you know, these are negotiated prices typically between the vendors and the hospitals and, you know, perhaps that bargaining parameters within that negotiation are different; information folks may be something that's driving differences in those

175

negotiated outcomes.

And then also you could think of things along the lines of just contract structure, right? So are there nonlinear contracts in here, some sort of bundling, some sort of exclusionary behavior, and so on?

And, finally, you know, you look at some of these much more commoditized products, and you think about, you know, what is potentially driving market power here? And it reminds you immediately of kind of the Stigler, you know, thinking about why is there price dispersion among buyers for a commodity product? Well, some sort of search costs, or think about here much more broadly, as I'll say search costs many times during this presentation, but what I really want you to think of is kind of the full set of things that go into forming a buyer-supplier relationship, right?

So think not only kind of finding a potential supplier, but all the kind of due diligence and work that goes into figuring out and developing a contracting relationship. And, you know, along those lines, you might think that these could be important in this particular setting.

So in this paper today, we're going to try and say something about all these aspects. So today the

176

distribution cost one is on the agenda but kind of not going to be underlying any of the things you see today, so just keep that caveat in mind.

On the contract structure, it's not going to be built into the model I'm going to show you, but we do a lot of work both in this paper and in a previous paper -- you can look at the 60-page appendices if you want to -- to just kind of do everything we can check, both in terms of, you know, qualitative efforts and interviewing people and also all of the checks that we could think of in the data.

Really, we find very little evidence of kind of any underlying kind of complicated contract structure or bundles underlying this. So for today we're going to think of this as mostly being driven by some combination of potentially demand heterogeneity or brand preferences, heterogeneity in bargained outcomes, and the heterogeneity in these kind of search or contracting costs.

Now, you might ask me, you know, this log 1 price, everybody knows it's not supposed to be a law. You just gave me a bunch of reasons why it shouldn't hold, so why should I be so interested in this? Well, you know, you can't get this kind of price heterogeneity unless it's all driven by these kind of differences in distribution costs, for example, without having pretty large markups, right? And so when we talk about what's underlying price heterogeneity, we're really in part talking about what's underlying what are potentially some relatively large markups in this industry, right?

And, you know, related to that, you know, knowing the sources of these markups is going to be important as we think about, you know, what might be the potential remedies or what we might expect to come of various policies or kind of things that are happening out there in the economy, right? So one of the things that, you know, we all worry about or people think about in every industry is when is Amazon coming?

So there's been a lot of talk in the medical supply industry. Amazon hired the COO of Cardinal Health about a year ago and has been looking into this area. So you can imagine, you know, what would that kind of information or that kind of intermediary do in a world like this?

You know, another thing, you know, these bundled payments and moves towards physicians maybe internalizing more of the costs of the decisions they make, you might think is something that would affect,

for example, brand preferences and so on.

And then, finally, you know, we are going to be looking across quite a -- you know, these are all medical supplies, but it's quite a heterogenous and large group of different categories here, so you'll recognize the approach that we take as being a very traditional IO approach, but we are taking it across, you know, a pretty large number of product markets, and so, you know, hopefully building towards this idea of, you know, we want more evidence in more areas along the -- you know, the -- building towards what the macroeconomists want from the IO world.

Okay. So I'll talk a little bit more about this institutional context, give you an idea for what we're working with here, the data set that underlies this entire endeavor, and then how we try and break down these pieces of the price dispersion.

Okay. So we're thinking about hospitals contracting with suppliers in, say, a given product market. So here I've just shown you the kind of set for coronary stents, because it's very simple. There are three vendors. It's probably the most concentrated, I think, of any of the markets that we look at here. And, you know, a hospital is going to be thinking about potentially contracting with its

vendors, and that's typically the job of an administrator at a hospital, so they are the one who's in charge of negotiating these contracts, making sure that there is something on the shelf when a provider goes there and needs to get something done.

Now, we drew the box there around the providers as well because, you know, perhaps there's input from providers in this, right? So, in particular, with some of these things, like physician preference items, like a coronary stent, a surgeon is going to be pretty upset if he goes and he likes to use the Medtronic stent and it's not on the shelf when he goes to use it, and the administrator is likely to hear about that. So an important part of what we're going to do today is going to have to be thinking about the formation of these buyer-supplier, you know, choice sets, potentially with input from the people who are actually using them, the physicians. And then, finally, once things are on the shelf and there to be used, providers are going to make decisions as patients come in in order to do their best to treat those patients.

And then another important feature of what we're going to do today is that this sort of activity is going to happen across lots of different product

categories within the hospital, some of them being relatively far away from one another. So what do I mean by far away? I mean something like, you know, coronary stents, they are used in interventional cardiology in the catheter lab, versus, say, neurology devices, which may be sold by the same vendor but sold by an entirely different sales force to an entirely different set of physicians and surgeons. And we are going to try and leverage the fact that there may be some linkages between these, you know, not on the demand side, but on the kind of cost side, on the administrator front, in order to get some mileage in solving some of the challenges of identification in this setting.

Okay. So a key that makes this whole endeavor possible is there's a really cool data set on basically everything these hospitals purchase, so it's about 20 percent of U.S. hospitals over the course of six years. We see all the purchase orders they issue, so prices and quantity, at the -- kind of the -- you should think of it as, like, the stockkeeping unit level, so not only just the product that you know the manufacturer and the vendor of that, but also, like, the size of that product, and it's for lots and lots of different SKUs across lots and lots of different product categories on a monthly basis.

And, you know, there are many challenges in dealing with this data and kind of working it into a format that we would do traditional supply and demand analysis with, and for today's short presentation, I will refer you to the many, many appendices, especially in the previous paper but some in this paper, about those. I guess right now I just want to make the -- start to point out, the way I'll try to summarize this is -- so in the paper, the current version, we have 24 different kind of non-PPI categories and six different physician preference item or PPI categories. So as I said before, these preference items are going to be things like pacemakers, drug-eluting stents, hip and knee prostheses.

The non-PIPs are going to be a little bit more heterogenous. So in there you'll see things that sound pretty commoditized, like surgical gloves, like sutures, and, you know, trocars are something that's used in laparoscopy procedures, a fairly common item, to things that are starting to get maybe more closer to PPIs, like a bone nail, right, but something that you maybe think is -- you know, a bone nail would typically be used, for example, in a prosthetic

procedure, but it's not kind of, like, the core item that's being put in typically in said procedure, right?

And so today mostly I'll just refer you to these rows that say "average," which is like the average of all these results across those two different big categories, but I threw in six of the line items just to give you a sense here, and the paper has all -- has the results for every single category.

So what you see immediately is these non-PPIs are used more often, right, so these tend to be kind of more ubiquitously used items both in terms of the numbers of hospitals that use them and the frequency with which they are used, but they are lower priced items typically, right? So actually once you kind of multiply P times Q, the actual spend on these PPIs tends to be about double of that of the non-PPIs.

And you'll also see, you know, as we documented in that first figure I showed you, that the prices are quite different across hospitals for all of these different categories. So whenever I show you a summary statistic in this case -- so, for example, the price is there -- that's going to be the quantity-weighted mean across all of our observations

of a price within that category, and that coefficient of variation is going to be similar to one I described to you before. So take a given product at a given point in time, look across hospitals, calculate the coefficient of variation, and then to give you just one number here for a category, we are just going to quantity-weight that number across all the different products in that category, okay? But, again, as you can see here, the quantity weighting kind of lowers us a little bit, so some of those huge coefficients of variation were coming from less used products but still quite large, on average 13 percent.

The other thing that you'll see varies here is the kind of size of the potential choice set, so there's the script J here, right? So this is the set of products used across all hospitals at a given point in time in the data, on average, and then that -- comparing that to the script J with the h subscript, which would be the size of the choice set we observe for a given hospital on average in the data.

So you can see, you know, we're looking at something like 10 percent or less of all the products that all hospitals are using will be used on average by a given hospital, and there's a lot of variation in that measure as well across hospitals, even more than

there is in the prices, right? So you have some hospitals that source quite a few different things from different vendors, some hospitals who source only a few.

And then, you know, there are some other hints in there -- in here that there may be some combination of either contracting frictions or heterogeneity in preferences. So just a few kind of, you know, simple statistics that start to get at this is if you take J star here to be, say, the most commonly used product in a given category, how frequently is that most commonly used product in the choice set of a given hospital, all right? So about 34 percent of the time for the non-PPIs versus 60 percent of the time for the PPIs. And similarly, you know, how often is that actually also the most used product within a given hospital, right? So how it kind of correlated our purchasing patterns across hospitals, and only 16 and 25 percent of the time.

So, you know, it's -- we find it at least pretty striking that you kind of see all these hints of lots of heterogeneity in purchasing decisions, and, in particular, in some of these non-PPIs, where you might think, ex ante, at least, it's kind of our prior that there's kind of maybe less inherent

185

differentiation among some of these products, right?

Okay. So just a few kind of things to talk about, kind of what may and may not be underlying some of this variation we see across hospitals. So it turns out, looking at prices, there aren't too many observables you can throw at it that explain too much of the meaningful price variation, so -- you know, in terms of -- observables in terms of hospital characteristics, so this is looking at -- just for stents, looking across bed size bins of hospitals and box plots for each bed size bin. As you can see, kind of no real discernible pattern in terms of bigger or smaller hospitals getting better deals.

It's the same if you look at stents for other hospital characteristics, like is it a teaching hospital or not? Is it a public or private hospital? And, you know, if you look at the relationship between price and bed size across all these different product categories, for some there will be a -- you know, a negative relationship, for some a positive relationship, but invariably, it's a pretty small relationship. So it's never explaining a lot of the variation that we're seeing in prices.

Similarly, these choice sets, you know, getting back to, you know, the institutional setting that we

186

were talking about before, one of the things that is quite predictive, actually, of whether or not a product is in your choice set in a given category, in a given hospital, is the spend of that hospital with that same vendor in other hospital categories, and you can do this by other hospital categories that are near and far here. I'm just showing for far, because those are going to be the ones we are going to be interested in as kind of giving us some leverage here for identification, but, you know, if -- in some of these product categories -- so this is plotting coefficients across categories here in a regression on, you know, product time dummy variables, vendor HRR -- hospital referral region -- dummy variables, hospital fixed effects, and looking at, you know, what's the difference between someone who's above the median or below the median in terms of spend on these far away categories. You know, for some product categories, it's quite dramatic, being above the median, you know, like double your propensity to be in a given hospital.

And then, finally, how do these two things correlate, right? This is obviously a quite speculative regression that doesn't have, you know, a lot of causality behind it, but nevertheless, I think, you know, interesting to think about, you know, for

187

people who have larger or smaller choice sets that they're sourcing from, you know, how does that relate to the prices that they're paying?

And, you know, in particular we think of it as interesting in thinking, you know, back to these issues of, you know, one reason you might have a small set of suppliers would be, you know, you have these contracting frictions that keeps your set of suppliers smaller than they might otherwise be. Another reason would be you're doing this strategically. You're excluding some suppliers so that you can leverage better prices from the suppliers that you, in fact, do buy from, right?

And, you know, the prior would have this very strong, if this were like a well identified kind of causal regression, a very strong prediction of a negative relationship between the size of the number of people you buy from and the price, whereas, you know, that prediction would be a little bit more complicated in the second.

And so, you know, we do find it at least suggestive, the evidence from this, that the relationship tends to be -- tends to be negative between these two things, and we do, you know, a little bit more work on this both in kind of this more

188

reduced-form analysis and then in some ex post analyses after we get our demand estimates that I won't have time to go into today, but our take-away from the entire endeavor, that at least in these product categories, it seems that the story is not really one of exclusion being a strongly suggested thing that's going on in this data set.

Okay. So how are we going to, in fact, disentangle these features, right? So we kind of have these three items that I told you about, all interrelated with one another potentially, and we're going to have to think about how we disentangle them.

And so, you know, what we're going to do is think about, you know, a model where, you know, a hospital has some ex ante beliefs over the qualities of some products, so maybe some product time quality, some, you know, vendor HRR quality on average, my hospital kind of needs, you know, preferences on average, with some unknown components, at least unobserved to the econometrician, these XCs, right, and kind of the twist on what we're used to seeing here is going to be that we're going to have these XCs that are unobserved to us, but one of these, the XCO is potentially observed to the hospital before they go out and contract with a vendor. So this is what I was talking about before with this idea that, you know, maybe physicians are influencing administrators and making these sourcing decisions, and that might be moving things in a way that's difficult for us to see as a researcher.

You're also going to have some known idea over what marginal cost would be, some bargaining weights that, again, you'll kind of have a sense of, on average, what happens with a given provider, on average what happens with my hospital, but some realization of the joint bargaining split that is kind of yet to be discovered until I actually do my due diligence and kind of pay these costs to go and think about actually contracting with these vendors.

And then, you know, finally, these choice sets, so these script Js are going to be determined. You will learn these unobserved portions, and contract prices will be set. Those will be set for some period of time, and physicians will treat patients as they come in, and quantities will be realized. So the two challenges in this setting is -- one is kind of the traditional one that we're used to, is price may be some function of things we don't observe in the demand system, and then this kind of -- this different feature where your actual choice set might be a

function of your preferences if physicians are influencing the choice set.

Okay, and so our approach for the latter is going to be to look for items that are pushing around search costs and, therefore, pushing around the choice set. And our approach is going to be very similar to, you know, the traditional selection correction that we all learn kind of in the labor context, and, you know, in this case, you know, the preferences of hospitals might be -- who actually buy a given product might be higher than the average hospital out there, and we are going to use a control function approach where what we are going to do is estimate what's the expected value of this unobservable for a hospital that actually contracts for this given product, and it's going to be based on kind of a reduced-form version of what you might think of as a search model that's going to include these far away spend variables as the excluded instruments that are uncorrelated with demand but only correlated with search costs in the choice set.

Okay. Then we're also going to have to tackle our more standard price endogeneity problem, and there we're going to leverage the previous paper that we wrote with this data set. So the reason this data exists, it's a hospital benchmarking platform that

actually provides information to hospitals on what other hospitals are paying for prices or paying for different items, and in that paper we found that it seemed to be highly suggestive of an asymmetric information story, where when you found out that you were really in the far tail of prices for things that you were purchasing a lot of, your prices tended to go down subsequent to getting this benchmarking information, and we are going to use that here as an instrument that's shifting around price exogenously in order to help us get some identification on the price coefficient.

And so we're going to have a -- you know, a demand and supply system here that's going to be kind of a simple nested logit with a nest on the outside good, hospital fixed effects, product time fixed effects, you know, the selection correction, kind of your standard Heckman type thing with the demand instruments that I just mentioned, kind of standard Nash-in-Nash bargaining problem on the pricing side, where, you know, a bunch of those parameters are going to come from the demand side. We're going to parameterize marginal cost in bargaining, you know, embedding that sort of "do you have access to information or not" inside the bargaining

parameterization, and jointly estimate the whole thing via GMM.

All right. And so this would usually be the part where I would tell you about how we go about estimating search costs, but since I'm already actually out of time apparently, it's a good thing that I'm skipping ahead and going to summarize that for you briefly. So what comes out of the demand and bargaining estimates, the thing that was probably most surprising to us was this extreme price insensitivity across kind of all these product categories, right?

So PPIs are much less price-sensitive than non-PPIs, but all of these categories, you just don't see much price sensitivity in demand, and so that's a big thing that's underlying these kind of markups that we're seeing in this market. So the markups on average are, like, 20 to 80 percent, and so actually it turns out that the fact that prices are negotiated, that hospitals have some monopsony power is really a key element that's keeping prices down actually relative to what they otherwise would be in this market.

So just two exercises to wrap up that we looked at to try and disentangle more what's going on in this market, so what's the role of these search frictions.

The first we look at, you know, in this -- you can imagine if you're familiar with these various search models, what we're talking about here is a relatively complex problem, right? You're sort of searching for a set of suppliers who you're going to kind of continually purchase from, so basically a portfolio that you're searching over. There's lots of heterogeneity in the demand and pricing specifications that I showed you, and so this is going to be a very complex search problem with large potential state spaces.

So the approach we are going to use to try and actually estimate these search frictions is going to be using moment inequalities, you know, based on some necessary conditions for products being in the choice set, and I think the slightly -- you know, the slight innovation or twist we have on some of the other papers that have been out there in this space is we've come up with these kind of loose conditions that we argue are consistent potentially with many different models of search or choice set formation.

Those costs end up being about on the order of 10 percent of price, so, you know, meaningful when we think about what markups are out there, but not huge compared to, say, like the price insensitivity.

Then finally what we do is a decomposition exercise where we, one, shut down bargaining variation, see what kind of variation we see across hospitals in that case, recompete equilibria, and so on. Two, shut down demand estimation, see what kind of variation we see across hospitals in that world. And then maybe more interestingly just do a very extreme counterfactual where what if everyone -- there were no search frictions? Everyone had access to the entire choice set that's available out there, what would we see?

And what we find actually is that the prices would go down a little bit, but not a ton, right? So on the order of something like 5 percent price reductions you're seeing here, and what you're seeing -- you know, much bigger effects that would come from that is potential, you know, consumer surplus gains through the additional variety and access to quality, right?

I don't want to hang my hat on that totally, because those who have worked with these models know there's a lot of extra logit errors being thrown in there, in that consumer surplus analysis, but I think the take-away is here is that the price implications of contracting frictions don't seem to be huge here.

You know, the big story is this lack of price sensitivity in this market that generates a lot of market power, this bargaining by hospitals that holds that market power down, but there's a lot of heterogeneity in these bargaining parameters that we're estimating here that's leading to the price dispersion that we see.

So we still have plenty of work to do on this paper, and I'm way over time, so I will stop and listen to your comments on that. Thank you very much.

(Applause.)

MR. PETEK: So Tobias Salz will discuss Matt and Ashley's paper.

MR. SALZ: So, yeah, let me already start by saying or thanking the organizers and Ali for allowing me to discuss this paper. This is something that I'm personally very interested in, and it was a lot of fun to think about it, and I really like this idea of this decomposition. So if you work on search models, you are oftentimes asked, well, is this not just some sort of bargaining friction instead? And I would not disagree with that. So I think this is a super valuable exercise and -- yeah, so I'll jump into the details here.

Let me spend one slide actually on motivating

this question. So the -- this literature, I mean, very broadly -- and I may be oversimplifying at this point -- asked, what is reflected in a price? And this started with the observation that we see a lot of price dispersion in markets where we shouldn't expect it; namely, for example, retail financial products where we can arguably very well control for all the quality attributes that buyers should care about. And the explanation is that there's some sort of friction that prevents buyers from picking the optimal -- the optimal product for them.

And recently, there's also a literature that argues that we see a lack of state contingent pricing, so sort of the opposite of this, right, that -- in particular, within retail chains across locations, there seems to be less catering to local demands than we would expect. And one explanation that's been brought forward here is that this might be due to some sort of managerial fixed cost that prevents firms from charging the optimal price.

And so, arguably, in this paper, both of these explanations are relevant. So this is a business-to-business market where suppliers might, due to some frictions, not set the optimal prices that shows up here in the bargaining, in the relative bargaining parameters, and on the demand side, there is a friction that prevents from finding what the optimal set of products would be. And what makes life even harder here is that on top of these two explanations, there might also be, of course, preference heterogeneity and cost heterogeneity.

Now, as Matt has pointed out and as you are probably well aware, you know, what -- the true underlying reason for price heterogeneities in this market will, of course, very much determine how we want to think about policy. So if it's true that this is due to search frictions or some sort of informational frictions, then what has recently become popular in healthcare markets to provide information about prices might be very valuable. If instead this is due to preferences, then, of course, that would be a different story. If you think about mergers, then heterogeneity and bargaining ability are, of course, important to understand, okay? So I think, again, this is a -- it's a very valuable exercise.

It's also something that is actually -- it takes about two minutes to find a lot of corroborating evidence online from the view of practitioners. So this is something that's very much in the minds of practitioners, so this is -- these are two quotes from

a website that's called Healthcare Finance, where they basically describe that it's very important who the person is that you pick for these negotiations and that different hospitals have different ability to solve this problem, and it's informationally a very daunting task to, you know, keep track of all the prices, all the vendors, and the various ways in which you could purchase these things, okay? So there's definitely a lot of supporting evidence for what the authors have in mind here.

Okay, before I jump into specific comments on the model, let me quickly recap -- and I'm actually happy that I'm recapping, because I think Matt did not get a chance to go over the search cost estimation, so I'll hopefully cover this.

So this is a model where hospitals have preferences over items that they want to source, and then at some -- in some costly process, they can add those items to their consideration set. And then there is Nash-in-Nash bargaining, so this is a standard Nash-in-Nash bargaining framework, but within the set of items that have been added to this consideration set, okay?

And so this is basically the sequence of events, and now the estimation goes in reverse order

of how these events occur, so they estimate jointly a demand and bargaining model, where they keep track of selection into this consideration set by using this control function approach, and then in the last step they get at these search cost parameters, and the idea here is that they resolve this bargaining game to get a new set of prices from which they compute the added inclusive value of adding a specific item to your consideration set, okay?

And then they basically, from those added inclusive values, get conservative bounds on the search costs. So you get a conservative upper bound by saying that a product that we see in your consideration set must have been added at some point, and the most conservative bound is by adding it to the empty set, right? That's when it's providing the highest value.

And conversely, a product that is not in your consideration set would provide the lowest -- so a conservative lower bound -- the lowest value if you add it to the entire set of products that's available. That's when it provides the smallest marginal value, okay? So that's how they, in a parsimonious way, without having to take a stance on how the exact search cost model looks like get these bounds.

Now, what you see, though, is that all of this depends very heavily on getting right the consideration set, right? And so that's unfortunately not something that's directly observed here, and I think the authors made a very sensible assumption in saying that this is the set of products that have -- you've seen purchased in the past, right, so that's natural in that it leverages the panel structure that the authors have access to, but -- because it plays such a crucial role in the identification of the bargaining parameters, but also on these bounds, I want to push a little bit here.

So one simple thing one could do is to simply make a -- instead a rolling window assumption and look -- you know, sort of varying the length of this window and see how robust the results are to different assumptions here. But pushing this a little bit further, what could also be exploited is the fact that consideration sets lead to specific asymmetric substitution patterns. So we all know that, you know, we get other ways of -- we get asymmetric substitution patterns in other ways in demand models, but consideration sets say that there are asymmetric substitution patterns along the boundary of the consideration set, right?

And so I'm wondering whether the authors can use this insight, which has recently been formalized in Abaluck and Adams, and I'm sure there are other papers that I'm not aware of, and so my guess is that their approach is a bit restrictive for this market with business-to-business and contract-specific prices, but one thing that they maybe could do is to take their definition of this consideration set and see whether something that according to this instrument that they use for the control function approach gets randomly placed in the consideration set, has other substitution patterns than something that's outside of the consideration set.

Now, the problem with this is that this is not a posted price market, right? So you cannot just look at price variation and sort of see how it -- how substitution patterns adjust, because every business has a specific price that depends on the relative bargaining parameters and other attributes. So my suggestion here would be to maybe use this benchmarking database and treat it as a posted price, okay, and see whether, with that, you get -- you can test for these asymmetric substitution patterns that you would expect to see if you get this consideration set right, and maybe you can also test what the most

likely consideration set would be. And I think that that would sort of go a long way.

My other comment is that at the end, the outcome of the paper is a decomposition exercise into preferences, relative bargaining strength, and search costs, and something that I'm wondering here is to what extent this might be driven by a very specific parameterization of these three different channels, right? So if we are looking at this -- we have, for example, this information variable in relative bargaining strength but not in search cost, and we have vendor fixed effects in the preferences but not in the relative bargaining strength.

So what I would like to know here is either, a priori, you know, do we have strong reason to expect that, you know, we have to put these objects there and not somewhere else, or, you know, do we want to be sort of completely agnostic and put all these things into all these three different types of channels, at which point, of course, you would pretty heavily rely on function or form assumptions, but, you know, sort of if you really want to get this decomposition right, I think you need some justification for, you know, why these things show up at these specific places.

So I'm almost running out of time. Let me just

make one more comment on something that I find personally very interesting. Actually, I learned about this when I was visiting here at the FTC a while ago, and Matt can take this comment with free disposal because it's really speculative, but what's interesting about these markets is that we have these group purchasing organizations here, which essentially every hospital is participating, so more than 95 percent of hospitals are part of these GPOs, more than 80 percent of all purchases are conducted through a GPO, and what they essentially do is they -- I mean, supposedly, you know, strengthening the bargaining power of hospitals, and also provide information about sets of products that are out there.

It would be interesting whether this can be, at least in a reduced-form way, be picked up by these bargaining estimates. I know it might be hard to get data on this, so that's why this is quite speculative, but, you know, these are sort of a fascinating entity that would provide some separate variation in information and bargaining strength, separate from preferences. So I think this might be quite interesting to study.

With this, I want to wrap up and say I think this is a really interesting and insightful paper.

I've been frequently asked about this, you know, decomposition to search and bargaining costs. I think this will be very valuable to the profession, and I think this is also something that brings together two different literatures, so the search cost literature has formerly been a bit separated from these vertical models, and so I think in that regard it's also a very nice paper.

That's all I have to say.

(Applause.)

MR. PETEK: So we have about five minutes for questions.

MS. JIN: It's a very interesting and sophisticated paper. I'm wondering, to the extent that both hospitals and suppliers are sort of long-term players, they know they're going to engage with each other for a long time, so to what extent do you see dynamic concern show up in the data?

MR. GRENNAN: Yeah, I think dynamics are a great -- like an interesting point here. It's something that we've incorporated a little bit by kind of just putting lags, like who were you with last time, and it's something that's kind of on our radar screen. Basically, the kind of two main things that we're still -- three main things, I guess, we're still trying to wrap our heads around how to do in a way that's tractable is kind of these dynamic issues, the fact that you have vendors selling many things both within a category and across categories, and this issue of kind of potential costs or returns to scale and distribution costs. So agreed, and we're in the market for very tractable solutions to these things.

MR. BRUESTLE: Hi. Steven Bruestle, Federal Maritime Commission.

If a hospital is buying a lot from the same vendor, could we possibly be seeing evidence of bundling, maybe I'll give you a cut on this product in exchange for you paying a little more on that product?

MR. GRENNAN: Yeah, absolutely. So this is -- our prior going into this project is that we were going to be doing a lot more on trying to, you know, figure out how to extract information on kind of unobserved bundling and contract features that we weren't seeing but that were probably there. We were kind of surprised that in the first paper, when we went to go do a lot of qualitative work in just talking to people, how infrequently they seemed to talk about these things being an issue, and then once we actually went to the data and kind of the set of tests that we've thrown at it in terms of correlations

between prices that you might expect to be part of a bundle, for example, or correlations in changes in prices or the comovement of prices, we're just not seeing much there.

MR. BRUESTLE: Okay, great. Thank you.

MR. BESANKO: So I wanted to build on the last comment that Tobias made about the bargaining weights. I thought that was actually something -- that was something that really caught my eye. As I recall, you said the bargaining -- the estimates of the bargaining weights for the vendors are somewhere between 1 percent and 42 percent.

MR. GRENNAN: Yeah.

MR. BESANKO: So there's a lot of bargaining power by the hospitals. Do you know anything about -- you know, are they larger hospitals? Are they hospitals -- are they hospitals -- or are these categories where there's a lot of bargaining weight from the hospital more commodified, more vendors? I mean, what can you tell us about the circumstances under which those bargaining weights differ?

MR. GRENNAN: Yeah, thank you for reminding me, because I left out the other thing that's on the agenda that perhaps our RA is sitting in Philly running today, is actually bargaining weights on

stuff. So this having hospital stuff to run it on is a relatively new thing that we have been able to do in a de-identified sort of way and something that we're very curious about.

I would say anecdotally, just in our conversations with people, like, there didn't seem to be a lot of correlations based on what we would have thought, ex ante, in talking to hospital purchasing professionals, where people who seem to be good at this were going to be, right? Like, it seems to be very person-specific, organization-specific, probably variables that we're probably not going to be capturing in, like, things in the AHA or that we're seeing here.

MR. RASMUSEN: If I could follow up, this is making me think of the Piketty Saez paper on CEO pay and market capitalization, because you're saying that getting a good purchasing guy is really important. If we could look at their salaries, for example, we'd expect those to be higher in the bigger hospitals, but maybe some smaller hospital thinks it's getting a real whiz at bargaining.

MR. GRENNAN: I mean, that would be interesting. Like, we constantly, in having these conversations, you know, do you get paid when you get

a better -- you know, when you're getting better deals, you know, there does not seem to be any formalized structure for this. It seems -- despite the great quotes that Tobias threw up there, this does not seem to be a super-mature market, as far as we can tell, in terms of, like, hospital purchasing expertise.

That doesn't mean that in some places it's not a big deal, but I just think it's something that there's a lot of money being left on the table through some combination of these, like, managerial fixed costs and incentive issues and professionalization of an industry and some interaction between those.

MR. RASMUSEN: Actually, you wouldn't want to use a high-powered scheme, because if you have a guy this tricky and good, he could really scam you if you gave him a percentage of amount saved or something, but it would show up in flat salary, I think.

MR. GRENNAN: We should look. No, I mean, to the extent that we can --

MR. RASMUSEN: Maybe we could get top five salaries.

MR. GRENNAN: -- get anything that proxies for that, we should try and think about that. Absolutely.

All right. Thank you very much.

(Applause.)

(End of session.)

PAPER SESSION:

COMPETITION, ASYMMETRIC INFORMATION AND THE ANNUITY PUZZLE: EVIDENCE FROM A GOVERNMENT-RUN EXCHANGE IN CHILE

MR. PETEK: All right. Our next speaker is Gaston Illanes, who's going to present Competition, Asymmetric Information, and the Annuity Puzzle: Evidence From a Government-Run Exchange in Chile.

MR. ILLANES: So, hi, everyone. Thanks a lot for having me. I'm very excited to be here. This is joint work with Manisha Padi, who is at the University of Chicago Law School.

So there's a vast literature in public finance documenting what is called the annuitization puzzle. This is the notion that, despite theoretical models predicting that retirees should allocate a large percentage of their wealth into annuities, in many practice, when you look at the outcome of annuity markets in the developed world, you see the opposite result. You see very low annuitization rates. If you look at the prices, annuity prices seem particularly high.

So the typical culprit for this outcome is adverse selection leading to market unraveling. Chile provides a really interesting counterpoint to this experience. In Chile, around 70 percent of eligible

retirees voluntarily decide to annuitize. Moreover, when you calculate the annuity prices at which they're annuitizing, they seem rather good. The markup over the actuarially fair annuity is quite low. So the broad, overarching question that we're trying to answer today is what lessons can we learn about this well functioning market that we can then apply throughout the rest of the world?

So how are we going to do this? We are going to build and estimate a really flexible, I think, structural model of demand for retirement assets. Our goal is going to be to recover the distribution of the underlying primitives that govern annuitization and welfare in this setting.

With those distributions, we are going to do two things. The first thing we're going to do is we're going to change the rules of the system to make the rules of the system in Chile look more like the United States. We are going to evaluate what happens to the annuity demand function and to the average cost curve and, ultimately, to the annuity market equilibrium when you move the rules of Chile to the rules of the United States.

As a preview, I'm going to show you that with Chilean preferences and Chilean rules, you get an

equilibrium that is quite similar to the observed equilibrium in Chile. With the Chilean preferences and the U.S. rules, you actually do get the U.S. equilibrium of the full market unraveling, okay? That's where we're going to go. Also, we are going to compute welfare changes, and we are going to try to compare welfare in both of these systems.

So the main take-aways that I want you guys to have from this paper is, first, we are going to find significantly more unobserved heterogeneity in the type -- in the preferences for these retirement products and significant correlation across the different dimensions of this unobserved heterogeneity than what has been posited by the previous literature.

Partly because of this, we can show actually that when you reform the Chilean system to make it more like the United States, you get something that the previous literature hasn't been able to get, which is full annuitization or close -- sorry, high annuitization in Chile and full market unraveling in the United States.

Having said that, the welfare implications are ambiguous. It is not clear. So, in particular, even though we can show that in the U.S. equilibrium, you get market unraveling, it is not the case that the

Chilean equilibrium Pareto dominates the United States equilibrium. There are people who are going to prefer the United States system, and there are people who are going to prefer the Chilean system.

Surprisingly, what we are going to find is that individuals who have a low value for annuitization prefer Chile to the United States, and individuals who have high values for annuitization prefer the United States to Chile, even though in the United States we can have market unraveling. The reason for that is because Social Security interacts with this market in a very specific way but drives welfare, and I am going to come back to that with more precision later on in the presentation.

So I need to teach you a little bit about the Chilean retirement system for anything that I'm going to do now to make sense. I will try to be brief. So Chileans save throughout their lives in private retirement accounts. You may have seen many people in this room, including myself, writing papers on this savings market. That is not the market that we're going to be studying today. Today we're going to study the market of what happens once you retire and you decide you want to access your money.

So to access this money, you are required by

law to go through an exchange. This exchange is called SCOMP. The way it works is you go to an office and you give SCOMP information about yourself, your age, your gender, if you're married, the age and gender of your spouse, how much money you saved during your working life, and you tell SCOMP the types of annuity contracts you would like to hear offers for. I'll be more precise about what an annuity contract type is in the next slide.

With this information and only this information, SCOMP collates everything and sends it to life insurance companies. Life insurance companies then decide, person by person, contract type by contract type, how much they are going to bid, okay? That information gets sent back to SCOMP. SCOMP ranks offers contract type by contract type, collates the information, and sends it to retirees, who then decide what they want to do.

The alternative to annuitization in this system is an asset called program withdrawal. Program withdrawal is basically a scheduled cake-eating problem that is frontloaded relative to an annuity, so you get more money right after you retire relative to an annuity payout.

The second crucial difference between program

withdrawal and an annuity is that under program withdrawal, whatever money is remaining in your account when you die is left to your heirs. So you can immediately see where adverse selection is going to come into this market.

If you are a 60-year-old, you have cancer, you have a high probability of dying within the next ten years, and you care about leaving money to your heirs, you're just going to put your money in program withdrawal, you are going to eat it until you die, and your heirs will get the remaining. On the other hand, if you expect to be long-lived, you have the incentive to annuitize.

So I mentioned annuity contract types. Annuity contracts here in Chile are quite sophisticated. They can have deferral periods, meaning that we contract today but they don't start paying out until d years in the future. They can have guarantee periods, meaning that we contract today, and if I die before the guarantee period is over, the contract continues paying out to my heirs. They can have up-front lump sum amounts, they can have step functions, and actually, you can mix everything I've said together. So contracts can become quite, quite complicated.

So what are we going to be working with? We

have an administrative data set of every single individual who has retired in Chile between 2004 and 2013. We have everything life insurance companies see about retirees and more; particularly, for example, we know in which municipality they live, which life insurance companies do not know. We see every offer that is made in the system. We see every choice that is being made. This is over 230,000 retirees and over 30 million annuity offers.

Moreover, we have been able to match this data set to the administrative death records. So we are able to tell, by 2015, whether these people are alive or dead. And for the purposes of this talk, I am going to focus on single life annuitants. If you're interested in why we did that, we can talk about it offline.

So there's a lot of descriptive work in the paper which unfortunately I don't have the time to talk about. I do want to hit the highlights, because I think they set the stage for what we're going to do next.

So, first, the market is very, very unconcentrated. There's roughly 15 life insurance companies making bids on people at any time. HHIs are very, very, very low. As a result, markups are partially -- as a result, when you compare the annuity markups to the actuarial fair annuity, they're very, very markups. Annuities are very competitively priced.

There's vast heterogeneity in accepted contract types. So it's not the case that people cluster on taking one particular annuity contract versus the others. I have motivated why this is. When there's heterogenous preferences, there's going to be people who are going to prefer contracts, for example, with guarantee periods, because they expect to die and they care about leaving money to their heirs, so on and so forth.

Markups are low. There's adverse selection into new annuitization. We can run the standard Chiappori and Salani reverse selection test, and we find what you would expect. People who buy annuities live longer.

And in terms of exertion of market power or in terms of exertion of brand preferences, roughly 20 percent of the population take what we call a dominated offer. By that I mean they accept an offer when there is another offer on the table that is the same contract type and is more generous from a company that has equal or better risk rating. Despite the

acceptance of dominated offers, the money people leave on the table when they accept a dominated offer is very, very, very low.

Okay. So what we're going to do in this paper is we're interested in making comparisons across contracts that have vastly different time properties in terms of being financial assets. So to begin, they have different flow payments over time. Second, they have different exposures to risk, both longevity risk and bankruptcy risk. And third, they have different inheritance properties.

So the way we're going to do these comparisons across these contracts, it's just really simple, and it's to set up a finite horizon consumption savings model. So the model is going to have the following features: We are going to make a model that has uncertainty over your own longevity and uncertainty about whether the company that you are annuitizing with is going to go bankrupt or not. It's going to have a CRRA utility function to allow for the possibility of risk aversion, and it's going to have the potential for a bequest motive, and by that I mean the potential for individuals to receive utility out of leaving money after their death so that their heirs can consume it, okay?

With this model, given a level of risk aversion, given a level of wealth outside the system, given a level of bequest motive, and given an expectation about my own mortality, if I give you an annuity contract offer or if I give you a program withdrawal contract offer, I can calculate the optimal consumption savings problem, I can solve the optimal consumption savings problem, and I can recover the value of that annuity contract.

The way to do that is numerically through the endogenous grid method or the grid points method. Sorry. So from now on I'm going to call a combination of risk aversion, outside wealth, bequest motive, and mortality shifter a type. And what we're going to do in order to estimate demand is to take a grid over this type space and solve the optimal consumption savings problem for every point in the grid, for every one of the 1.2 million offers that we see, okay?

Given a type and given a person, we are going to impose or we are going to assume that the individual accepts the offer that gives them the highest utility from the optimal consumption savings problem, and with that assumption, we are going to solve for the distributions of types that rationalize choice.

In this slide right here, then the problem of solving for that distribution of types is in the second set of equations. You can see that it's actually a simple minimization of a constrained OLS problem. Pi here is the probability that every single -- associated to every single type. This is a PMF. It must sum to one, and each of the elements must have non-negative probability associated to them.

This may look familiar to you because this is just Fox, Kim, Ryan, and Bajari. The only contribution we have here is that we're marrying the Fox, Kim, Ryan, and Bajari framework to an optimal consumption savings model. Yeah.

So you may have some concerns about this model. I'll point out the ones that I have. To begin, it's a purely financial model, and what I mean by this is that people are going to accept the offer that gives them the highest utility. As a result, there is no scope for brand preferences. One of my advisors was fond of calling this the Snoopy effect because one of the companies in the system was Met Life. So the idea was perhaps you like Snoopy and, as a result, you are willing to accept a lower offer from Met Life than you would from another company just because you like the brand. We are ruling that out. I'm comfortable

221

ruling it out, to be honest with you, because even when we see the acceptance of dominated offers, the amount of money that is being left on the table is rather low, but that is an assumption.

Second, there could be information revelation in the request stage, and by that I mean when you elicit contract offers, the contract menu that you are requesting could tell insurance companies information about your own immortality. If that is the case, we are ruling it out. It would bake in correlation between the choice set and your own distribution of types, similar to what Matt talked about in the previous presentation.

To alleviate that concern, we're working on re-estimating the model conditional on the request set so that within the request set there is no heterogeneity and no information revelation. The hairy thing here is going to be finding a group, a mass of consumers, that all request exactly the same contract so we can run this.

You may think that there is heterogeneity in distribution of types across observables; for example, it might seem insane to estimate this model jointly for men and women. We agree. We're separating out across genders, and we're also separating out across

222

pension savings quartiles. So we're going to estimate this model for every gender/pension savings quartile pair separately.

And second and finally, those of you who have worked with these types of estimators may have experience that they can be quite finicky and sensitive in terms of the grid that you are choosing. We're trying to be very careful about the choice of grid and trying to pick it in a smart way so that this is robust. I, unfortunately, don't have the time to delve into that, but I'm happy to talk about it offline with you if you are concerned about that.

So a key question that you might be thinking now is, how can you identify these distribution of types just using the choice data? And from a formal perspective, what you need is -- in the previous slide, we had this S matrix, which is simply a matrix that has, in every row, individuals and offers, and in every column, it has types. This S matrix is going to have zeros and ones, a one when a type chooses a contract and a zero when a type does not choose a contract, and formally what you need for identification is invertibility of S-prime-S. Now, what does that mean in practice? It means that different types have to make different choices.

223

That's the only way we are going to be able to recover the distribution of types.

Let me give you an example of when that breaks down. Risk-neutral individuals do not choose over lotteries taking into account their outside wealth. So for the risk-neutral types, we, of course, cannot recover the distribution of outside wealth. Despite that, I think that this works rather well, in particular because the choices here that people have over these different -- what you could think about as lotteries -- are quite stark.

For example, an individual who is illiquid upon retirement and who expects to live for a very short time will never take a deferred contract even if the deferred contract is quite generous just because they won't live long enough to recoup the investment of not being paid for a certain number of years.

As another example, someone who cares absolutely nothing about leaving money to their heirs will never take a contract with a guarantee period, because a guarantee period only shifts down the payments you get over your life at the benefit of leaving money to your heirs.

Okay. So the unfortunate thing about these grid estimators is that the result of the estimation

224

routine is a list of types with different weights, which makes it hard for presenting. The list of types and weights for every single quartile gender is in the paper. I'm just going to talk about the highlights.

So the first thing that we found very interesting is that there's a large, significant heterogeneity in bequest motive -- there's actually bimodality in bequest motive -- and that an intuitive result, we're finding bequest motives are higher for women than for men. This is consistent with findings in the development economics literature as well.

We're finding a large heterogeneity in mortality expectations relative to the table, that's the Chilean death table; that is, individuals are not discounting the future as if they expect to die according to the Chilean death table. There's people who expect to be sicker and there's people who expect to be healthier than the Chilean death table. Poorer individuals across the board seem to have higher mortality probabilities.

We're finding that the distribution of outside wealth that we are backing out shifts to the right as pension balances increase. We're finding low heterogeneity in risk aversion, significantly lower values than the literature, and we're finding

225

mortality probabilities that are negatively correlated with bequest motives and that are negatively correlated with risk aversion. This is really important.

In a standard adverse selection market where the only source of private information is just mortality, the first people who annuitize are going to be the people who expect to be the longest lived. The last people to annuitize are going to be the people who expect to be the shortest lived. That creates the standard increasing average cost curve result.

Here, it doesn't have to go that way. It could be the case that the first people who annuitize actually aren't the people who are the longest lived, and I'll show you how that happens and when that happens.

So the remainder of the talk, I am going to start actually applying these results. So the first thing I am going to do is I am going to simulate market equilibria under stripped-down, simple versions of the Chilean and the U.S. institutional framework. My goal here is going to be to highlight the change in the demand in the actual cost curve that's induced by the introduction of Social Security.

In both Chile and in the U.S. -- and in

226

everything you are going to see now, I am going to assume that there is a single annuity contract, zero guarantee, zero deferral period, and I am going to assume that the market is perfectly competitive, and I am going to assume, just like it is the case in Chile, that pricing is on gender and on pension balances.

I'm going to allow for the possibility of fractional annuitization, and by that I mean that individuals don't have to allocate their full wealth to either an annuity or to the alternative but, rather, they can allocate fractions of the wealth to both retirement assets.

And I'm going to assume that there's a 1 percent bankruptcy probability in the world where you annuitize. This is mostly to bake into the model the feature of the United States system where we take a private annuity and the company goes bankrupt, you're out of luck. The results that you're going to see now actually don't change if you change from 1 percent to 0 percent.

In Chile, the alternative to annuitization will be this program withdrawal problem that I told you before. In the United States, I'm going to follow Mitchell Perturba (phonetic) and co-authors in assuming that 50 percent of your pension savings are

227

taken away from you and returned to you immediately in an actuarially fair annuity. It's just that you have no choice over this matter. The remainder of your money can be either allocated into an annuity or it can be withdrawn lump sum, okay?

So with that we can start looking at the equilibrium for both Chile and the United States. I'm going to show you results for females in the second quartile, because it's actually the results that are the most stark. You can see the other genders and the other quartiles in the data. The main conclusions we're going to see at the back end of the paper actually are not going to matter at all.

Okay. So the green line here is the demand function. The red line here is the average cost curve. Why is demand upward sloping? This is just the standard annuity thing. On the X axis, I have the wealth annuitized. On the Y axis, I have the generosity of the annuity. As the annuity gets more and more and more and more generous, of course, more people are going to annuitize. That's why the shape looks like that, okay?

The average cost curve, you can just think about it very simply as the highest annuity offer that a company can make given the annuitant population and

228

still break even, okay? In a world where the only source of selection into annuitization is heterogeneity and mortality, the first people to annuitize are going to be the longest lived; the last people who annuitize are going to be the shortest lived. As a result, the offer you can make and still break even is going to be increasing as a function of the amount annuitized.

Here you see, in fact, that for some regions of our actual cost curve, the curve is, in fact, decreasing, not increasing, suggesting advantageous selection. Despite that, when you compare the equilibrium here, represented by the blue dot, the annuity rate that you see in equilibrium is, in fact, lower than the actuarially fair annuity. So the advantageous selection is just loke (phonetic), okay? And we're getting an annuity rate in the simplified version of Chile of roughly 55 percent annuitization. I should apologize and say that nothing here has standard errors. We're working on those, and my apologies for that.

Here's the U.S. equilibrium. So, again, the green line, demand, the red line, average cost. There is no intersection. Full market unraveling. To be honest with you, once you add standard errors, probably there will be an intersection between zero and 10 percent annuitization. Materially, the conclusions that we are going to reach are not going to change.

So you can see that there is a large contraction and rotation of the demand curve when you introduce 50 percent Social Security. Why is that? Because now, very intuitively, every single person in the market already has half their wealth in an annuity. As a result, the willingness to pay for the marginal annuity dollar, of course, has to fall. That's the contraction in the demand curve.

The rotation in the demand curve comes from a homogenization of risk across individuals induced by setting such a high floor. Actually, the average cost curve doesn't change that much. I'm happy to talk about that offline. So here we get full market unraveling.

Okay. Now, 50 percent is just a number that Jim and Olivia picked. You could play around with other numbers and see whether this result is robust or not. So in this plot, I am showing you on the Y axis the fraction of wealth that is annuitizing when you move the amount of money in Social Security from 0 percent in Social Security, where the only difference

between Chile and the U.S. is lump sum versus program withdrawal, and 90 percent of your money in Social Security.

So you can see that for around 50 percent of your money in Social Security and above, you are getting the market unraveling result. For values of money in Social Security below that, that is not the case, okay?

Now, up to now, I've tried to make no statements about welfare. You may be thinking that market unraveling should have an adverse welfare effect, in particular for people who value annuitization. In fact, we're finding that the story is not as simple as that. So we've calculated type by type and amount in Social Security by amount in Social Security the compensating variation that would leave an individual indifferent between being in the United States and being in Chile. Positive numbers here are people who have to be paid in the United States to be indifferent between the United States and Chile. Negative numbers are the converse, okay?

So the main take-away from these plots is that in none of these cases it is true that one system Pareto dominates the other, okay? There are always going to be people who prefer the Chilean system, and

there are always going to be people who prefer the United States system.

I have a minute. I'd like to characterize these types, so I'll be brief about that. What we find is that individuals who fully take up program withdrawal in Chile dislike the United States system. We are going to call these people people who have low values for annuitization. The reason is quite intuitive. These people are being forced to annuitize a significant portion of their wealth even though, for example, they're going to die two years from now and they really care about leaving money to their heirs.

They do not enjoy the benefits of the Social Security annuity, and as a result, when you move to Chile and you let them put their money in an asset where, upon death, their heirs are going to get something, of course, their welfare is going to be higher.

On the other hand, people who greatly value annuitization systematically prefer the United States to Chile, and this was surprising to us because we expected that when the market unravelled, that wasn't going to be the case. The reason why this happens is actually simple. For levels of Social Security where the market doesn't unravel, putting all your money in

an annuity in the United States has a higher return than putting all your money in an annuity in Chile, so, of course, these people prefer the United States to Chile.

When Social Security is so high that the private annuity market unravels, well, Social Security is so high that you're already getting the Social Security annuity for a vast portion of your wealth. The remaining dollars are the dollars that you cannot annuitize, and for those marginal dollars, the difference between annuitization and lump sum withdrawal is not as large as for the inframarginal dollars. As a result, even though their welfare does decrease relative to cases where Social Security has lower coverage, in fact, the United States for these types still dominates Chile.

Okay. So we've estimated this model of demand. We've started playing around with the institutional setup. The key take-aways that I want you to come up with is that, like predicted, when you introduce Social Security, you are going to get a contraction in the rotation of the annuity demand function. We do get market unraveling, like you see in the developed world.

Despite this, the Chilean system does not dominate, and in particular, high-value annuitization types are not worse off in the United States. The people who actually prefer Chile are people who we think haven't been thought about too much. It's people who reach 60 or 65 years old, and they're sick, they're going to die soon, and they're not going to enjoy the benefits of the Social Security annuity. We're starting to think about policies that potentially could benefit these types.

So, thanks.

(Applause.)

MR. PETEK: So J.F. will discuss Gaston and Manisha's paper.

MR. HOUDE: Okay. Thank you very much for having me to discuss this paper. Let me just start by saying this is a great paper, very ambitious, and this is actually a great example of a -- you know, a very tiptop IO paper, uses very good IO techniques to estimate, you know, of course, an empirical paper on a great empirical public finance question that we should care about. So it's not -- maybe not the, you know, perfect paper for this audience, but I think the paper is going to have a great future, because, you know, the question is really important.

Now, what is the paper doing? So let me just

briefly give you a short overview of the question. I am going to give you a very broad overview of what the paper is doing. So the paper estimates a life cycle model of consumption savings with adverse selection. I put in parentheses "advantageous selection" because, you know, it does play a role, because the model is rich. It does have, you know, a fairly rich model of correlation types, and though for some segments it does have advantageous selection and then applies it to the Chilean annuity retirement savings system.

Now, what's different? So I am really not a specialist in the U.S. security -- Social Security system, but what's different about Chile versus the U.S. is these two things. So, first, when you have to retire, you choose between -- essentially you're offered this menu, which is a competitive exchange, where you have these companies offering you these products, who are bidding for your contract, and this is actually a very competitive market, which, you know, as Gaston showed -- and I am going to give you the example in the next slide -- the prices are very competitive, and consumers have these menus that, you know, great for them.

And relative to the U.S., the other big difference -- and this is where the results come

from -- is that unlike the U.S., consumers have the option of moving all of their savings to this private security market. They also have a public option if they don't want to, so this PW option, but unlike the U.S., we do allow these retirees to use all their savings and put them in this type of annuity. So it increase the market size quite a bit, which essentially solves a lot of the adverse selection problem, and as a result, we do have a much higher takeup rate than the U.S.

Now, the richest question is, as Gaston put it, you know, what would happen if we subject the poor Chilean to the U.S. and will the market unravel, and the answer is mostly yes, and, of course, as I said, this is a really important question, because, you know, we are stuck with that problem here in the U.S. We do have the problem of how do we fund the Social Security system in the U.S., and this is a good step in answering that question.

Okay. So this is not a great scan from the paper, but this is what the -- you know, hopefully the people in Chile see better. So basically what you do see when you retire is you see this set of bids and -- well, first of all, there is two things that I was personally surprised, is, well, first of all, you have

quite a bit of competition, and these prices are individualized prices. There's, you know, full price discrimination, potentially, but there's not a lot of price dispersion.

Now, you don't see it here, but there's -- you know, the range of prices is very narrow. If you take out the outlier at the bottom, you know, the range of price is about 2 1/2 percent, at least in this table. I don't know how representative that table is, but, you know, more or less, you know, we're not far from the LIBOR price, you know, essentially.

And a part of that dispersion is explained by the riskiness of these life insurance companies, but, you know, more or less, you know, this is pretty much one price. And then price means the payment, and then if you take in the markup that these guys are receiving, if you take out the very rich and the very poor, you know, it's pretty much constant markup, okay?

So, you know, the paper talks a little bit -- so I thought I should include a little bit of that since this is an IO conference. You know, this paper talks a little bit about this has evidence of price discrimination. This is -- you know, maybe I was thinking about this weird, but this is -- you know,

237

this is not -- this is sort of the opposite of price discrimination in some sense because the poor are priced -- are getting a higher markup than the rich, and so we think that the poor have more price -- are more price-sensitive, but I think the reason this is not price discrimination is because the poor -- there's really no -- there's not a lot of competition for the poor in this market, and that's really what's going on at the bottom of the distribution, which actually, one thing in the next iteration of the paper, we might want to exclude these guys at the bottom, because it does generate a lot of random variation in prices, which might violate some of -- you know, when we talk about these guys taking up weird offers, you know, maybe that's coming from the bottom of the distribution. So that might be one thing that can explain this.

The other thing is that the paper mentioned that a little bit -- so, these guys sometimes use agents to shop. Sometimes they shop on their own. They also have an option of renegotiating these offers. And the paper says, well, they don't negotiate that much. They barely negotiated by 2 percent on average. Well, 2 percent is the range in this table, so that will eliminate more or less the

238

mispricing that we have when people take dominated offers. So it would be nice to know a little bit more what's happening with these agents and the renegotiation.

And the flip of the markup is also if you charge very high price, people are not going to take these offers, and so these low guys, these low wealthy guys are not going to take those offers. But the other thing that is weird is why is it that the wealthy guys are not taking these really good offers? So this is one thing I was a bit puzzled when I saw this table. Why is it that the wealthy guys were actually receiving offers with negative markup and not taking those offers?

Okay. So, again, I was not the right audience for understanding annuity markets in general, although I really cared about the question, so let me -- so it took me a little bit of time to understand why demand was upward sloping. I might have been very tired, that's also part of the problem, but, you know, at the end of -- and the paper is very clear in terms of, you know, how things work.

Now, what was going on here is that, you know, the governmental option, essentially the payments are decreasing, so people who are going to take the

239

governmental options is only people who expect to die very soon, okay? And so people who are going to take the annuity are people who expect to live very long. And from the point of view of the life insurance companies, these are the risky people, and so that's the adverse selection problem. So people who are buying the annuity are people who expect to live longer than their age suggests, okay? And that's the problem.

So you can -- basically the way Gaston construct this willingness-to-pay curve, because the model is actually -- is complicated, right? So it's not that trivial to figure out what is an indifferent consumer given the nonlinearity. So he's constructing this indifference point, you know, what is my riskiness such that my -- I'm indifferent between these two contracts, and then I can raise the price of the contract, and then I figure out what is my riskiness so that I'm indifferent. And as they raise the offer, I get different levels of riskiness.

And so as you go -- as you try to raise the contract, you get kind of people who expect to live shorter, less risky individuals. So, okay, so this -- again, so this is -- you know, I'm not in this literature. I was very surprised that there was a

240

corner of the econometrics literature that was thinking about upward sloping demand curve still, but anyway, that was just a small comment.

The U.S., though, so the U.S. -- I think this is -- you know, the most interesting part of the paper, is, you know, what -- how do we think about the U.S. Social Security system in this context, and it's very intuitive, and Gaston explained it very well, is that it is both the rotation and the contraction of the demand, because we're essentially insuring a lot of the risk by telling you, well, 50 percent of your savings is going to be automatically annuitized, and so there's less of a need to annuitize the remainder, and so people are willing to pay less. So there's just less demand for it.

Now, I was reading the draft and I couldn't figure out why the points were moving left, so that's a small point, but it would be nice, since there's actually quite a bit of advantageous selection in some of these segments, to understand, well, if there's advantageous selection in the U.S. market, well, how does this work in this market? So Gaston mentioned that the cost curve should not change too much, but in the simulation, it does change a little bit. So I couldn't really understand everything about that.

Now, let's talk about the structural model a little bit. So the structural model is more complicated in its pictures because it has four dimensions of heterogeneity. So it's not just my mortality risk. It's my risk aversion, it's how much I care about my kids, and the initial wealth. So there's four dimensions. All these dimensions are allowed to be correlated. So it's a very rich model.

So Gaston used this finite mixture approach to estimate that. So one thing that I was not clear -- and in the presentation it was more clear -- was how observed heterogeneity is accounted for, and so that's clear now because everything is estimated separately for male and female, but in practice, there is more observed heterogeneity than that.

Now, this is a quote from the paper. You know, what -- the one thing that -- and I'm more of a parameter guy. I kind of like the normal model, but, you know, when you look at the identification of these models, what's difficult is it is very black boxy, right? So the model is identified because it's identified, because the rent condition is satisfied, and so you lack a little bit of the link between the data or the reduced form and the parameters. What is -- you know, that's lacking a little bit. And

since these correlations are so important, it would be important if that would be sketched out a little bit.

So I have a comment on the next slide, but one thing that would help here is to maybe estimate a version of the model that would fit a little bit closer to the literature, like the Cohen and Einav type paper, that it uses parameteric model.

Now, finally, about -- this is an IO conference, so I have to talk about the endogenous prices. I do believe that market is competitive, but people do pay different prices, and they do accept rejected -- the dominated offers. So there's some room for endogeneity here.

So you talked about brand preferences, so that's one reason why that could be. Another reason is the fact that these offers are sometimes renegotiated. So think about the case of two guys in the deal who look up servicing equivalent who accepted different offers. Well, if I saw different offers in the data, the model is going to say, well, we have different unobserved taste, but it could be that the price is measured with error, because we renegotiated those prices. And so that might be one thing.

And that that could be, you know, one source of simultaneity, and so that would be one way of

correcting -- so I have one suggestion there, and this is where understanding a little bit better the role of the agent would help. And then the other suggestion was related to the identification. So one way of talking about identification a little bit is through these adverse selection tests. So the shepherd, the Chiappori and Salani is a test of adverse selection.

Well, you know, it is -- if you pass the test of adverse selection, it does tell you that there is an observed heterogeneity, so if you find advantageous selection, like we do in the paper, that means that for some consumers we should be able to find the opposite correlation, and so we should be able to find that in the reduced form as well. So there should be a tighter link between the structural model and the reduced form, and that would help in understanding the results.

Okay, I think I am out of time. Thank you very much.

(Applause.)

MR. PETEK: So we just have time for a couple questions.

MR. BESANKO: So my impression is that a lot of the countries that have private account systems have a minimum pension guarantee. So if you don't save

enough over your life, you're guaranteed a certain amount. Now, that sounds Social Security-like. I don't know if Chile has that, but if so, did you look at that, and how did that factor into your work?

MR. ILLANES: Yes, yes. So, thank you. This is something that I do talk about in the longer format presentation but I can't touch in 25 minutes.

In the background of everything here, there is a minimum pension guarantee. On the first slide, why I say 70 percent of eligible annuitants accept an annuity offer, it's because if you cannot fund an annuity offer that falls above the minimum pension guarantee, you are not eligible to annuitize in Chile. You are not in this market. You must take program withdrawal. Those people are not in my sample. That's why I have 230,000 retirees over eight years, which may seem like a small number to you.

So if you are poor enough, you are not in this market. Those are the eligible people. If you take an annuity, it must be above the minimum pension guarantee, so that's it. If you take program withdrawal, when your money falls sufficiently low such that you are below the minimum pension guarantee, the Government begins to top up program withdrawal, essentially subsidizing you.

That's in the model. It's in the estimation. We're controlling for that, and it's one of the observables that enters person by person into the national consumption savings model, but, yeah, it's something that I can't talk about in 25 minutes.

AUDIENCE MEMBER: So I have a couple of questions. The first one is, do you actually see a firm in your sample offering a menu of options to a consumer?

And then the second question is somewhat related. So are you worried that the consumer may actually misreport their savings balance so that they can actually get a better price?

MR. ILLANES: Thank you. So most firms bid for every single contract type, okay? So it tends to be the case that if a firm is bidding for you, it's bidding for you on every contract that you elicited offers for, okay? So from that perspective, the common thing is to see the menu.

Regarding misreporting, there can be no misreporting. The way this works is that the Centralized Exchange actually pulls the records from the savings period and sends them directly to the life insurance companies. So from that perspective, there is truthful reporting so that that can't happen.

Yeah.

MS. JIN: This is an interesting paper. You focus on individual choice of contracts. If we shifted gear to, say, the program designer, the Government in the U.S., at least, by offering Social Security, U.S. Government is functioning like an insurer here. So I wonder from that perspective what implication would your results have in terms of, say, the risk that Social Security Administration is taking in terms of insolvency versus kind of privatize all the Social Security money into individual accounts?

MR. ILLANES: Yeah, so let me touch on the only part of Social Security that our paper can talk about, which is what happens when you're taking money out. I don't want to talk about how people should feel about when they're putting money into Social Security and what they should think about when they're 20 years old, because that's not our paper.

From the perspective of what happens once you're 60 or once you're 65 and you're deciding to retire, you want to withdraw money, our main finding is that there are going to be people who are really going to dislike this contract, right? And what we're trying to work on now is to try to determine, if for those people we could offer them an alternative

contract that they would like more, that will not lead to higher expenses for the Social Security Administration.

And I can't comment on whether that's going to work out, but we suspect that if you can come up, for example, with a program withdrawal alternative that is sufficiently not attractive, so annuitization -- so people who like Social Security stay in Social Security, but it is sufficiently attractive for people who really dislike annuities to leave the Social Security annuity and convert it into program withdrawal, maybe there's a way to achieve that goal. But that's something that we're working on, and I don't know yet.

MR. PETEK: Okay, thank you.

(Applause.)

MR. PETEK: We will take a break until 4:30 and come back with Ali's keynote.

(End of session.)

KEYNOTE ADDRESS:

SEARCH, ASYMMETRIC INFORMATION, AND COMPETITION

MR. PETEK: All right, let's get started again.

All right, so Ali Hortaçsu is going to give our second keynote address, "Search, Asymmetric Information, and Competition." He is the Ralph and Mary and Otis Isham Professor of Economics at the University of Chicago. He's also a member of the American Academy of Arts & Sciences, a fellow of the Econometric Society, and a fellow of the National Bureau of Economic Research. His recent research has focused on industrial organization, auctions, search and matching models, production, and financial networks, with applications in finance, energy markets, and the internet.

Ali?

MR. HORTACSU: Thanks a lot, David, and thank you so much to the organizers for having me on the program to provide some input. Thanks again for, you know, putting together this program, and, you know, sort of in the first half of the program, I saw a lot of Stigler 1964, so this is more -- in the afternoon, we switched over to more Stigler 1961, to search models, so we're -- and it was great to see, you know, in the previous session, we had a lot about search,

249

but, you know, is it about search, about bargaining?
The second paper was a financial products market, and actually I was asking around. I'm sorry, I'm like woefully ignorant about where the jurisdiction of the FTC is, you know, do you guys work with financial products markets at all? You know, I'm not so sure, but I do think these are important markets to study.
You know, to study this, you know, this is probably sort of the longest list of co-authors I have written a paper with, but, you know, everybody contributed in very important ways -- except myself, we'll see -- and so this is joint work with Sumit Agrawal, John Grisby -- who is going to be a very promising job market candidate, maybe not this year but next year -- so my former colleague, Gregor Matvos, Amit Seru, and Vincent Yao.
So to talk about Stigler 1961, diagnostic for some sort of funniness business going on in the market is price dispersion, and mortgages seemed to fit that bill, at least from a prima facie point of view. This is a plot I have from a paper by Amit and Gregor and their co-author, Umit Gurun, on subprime mortgages. So they found, after residualizing a lot of demographic information, et cetera, on rates, mortgage

250

rates for the same type of contract, you know, you can see the X axis, the horizontal axis in terms of percentage points, big dispersion.
What we're going to start off with in this paper, in this project, is not the subprime but conforming mortgages. These are things, you know, basically Fannie and Freddie, sort of these government entities insure more plain, vanilla contracts with more risk, with higher credit borrowers, still a pretty large residualized dispersion, you know, people, you know, on order of percentage point, interquartile, and the distance here -- and it's not just in the U.S.
J.F., with Jason Allen and Rob Clark, has, you know, two very nice papers on the Canadian mortgage market that really inspired us to write this paper. They note large dispersions in the Canadian residential mortgage market as well, both residualized and nonresidualized.
So since 1961, right, so to explain this, one of the main drivers is information, you know, George Stigler puts it much better, sort of, you know, there's power in information, right? So if you -- really, consumers don't see all the prices, especially in, you know, a market like mortgages, where you can

251

see posted rates, but, you know, posted rates might not mean anything.
You have to put in an application, and they have to check your credit and see if -- you know, what rate you qualify for. So it takes a while. There's a cost to actually getting that quote from somebody, with, like, insurance as well.
And then you -- and then what causes this sort of price dispersion, Andrew in a comment earlier said, you know, the sophisticated naive decompositions, right? So the sophisticated consumers, you know, easy to search, they have information, but nonsophisticated consumers, they just don't know.
In some very nice models, like the Varium (phonetic) model, they just take the price or they do very little search. And this has become, you know, for better or worse a very attractive framework in consumer finance, precisely because in this type of market, you know, information is relatively hard to get, plus, you know, a huge area of products.
Gaston talked about the very large array of contracts that people have to sort through, and they're complex. You know, when you get the mortgage product, what you're buying is the contract you signed. I don't know how many in this audience who

252

got a contract actually read through all of those pages, you know, probably not, you know -- maybe a few, but -- and infrequent transactions, okay?
So -- and then -- and we talk about, you know, why people pay different prices, you know, who pays more, sophisticated -- who pays more? Are these unsophisticated people or sophisticated people, et cetera? And we have different -- and this leads to a lot of justifications for interventionism, especially from regulatory agencies, right? So we want to protect especially the vulnerable population of consumers from making bad choices, and, you know, we want to help prevent firms from exploiting naive consumers.
And a lot of these interventions are in the form of, you know, information treatments, in the sense of, you know, mandated disclosure of certain things, you know, putting prices up on, you know, web pages, you know, that everybody can access or, you know, plain sort of things like interest rate ceilings.
And just to preview, I have a very good student this year on the market to talk about interest rate ceilings, to advertise for him, you know, some policy interventions that go from very coarse to, you know,

rather subtle ones, okay?
But a lot of this intuition comes from our understanding of what I'm going to call standard product markets. Standard product markets, the buyer's payoff depends on the price. The seller's payoff depends on the price, right? But credit products are somewhat different, and Gaston's paper was on adverse selection in the annuities market. This is a market where, you know, the borrower cares about essentially only the price. They might care about Snoopy, but -- you know, probably they shouldn't -- but the lenders definitely depend on who signs the contract, right?
Is it -- you know, aside from the rate you get, right? So is this person going to repay the loan or is this person going to, you know, pay it back way too early, and is there a repayment risk in this thing? So the lenders are going to screen. They are going to get a lot of information, decide whether to accept or reject applicants, and they put a lot of resources into this.
And I would like to say sort of a main, you know, let's say thrust of this paper is to sort of motivate that the screening aspect is very important in these markets, and I would like to -- we haven't

done it yet, but follow it on with other work, you know, trying to get at the importance of screening technology in these markets, especially in these days with -- you know, when people talk about big data and how this interacts with it, I think it's an important question.
So to get at this, we have a very nice data set -- although I keep hearing from people that even nicer data sets are coming online, so we should hurry up and try to publish this paper before the young people, you know, publish their papers faster than we do.
So we get -- there's essentially two separate data sets. We have data on mortgage applications, so we have all the information they filled out on those application forms and the decision, whether this was accepted or rejected for a mortgage, and we also have data on granted mortgages, the mortgages in the marketplace.
So the main, I guess, fact that I'm going to say -- and there's a few different figures that are going to say the same picture -- but here higher search interest intensity is going to be correlated with higher interest rates that people get, conditional on many observables, you know, controlled

for. Of course, a standard product market model where there's no screening, you would expect people who are searching more to be the more sophisticated consumers, that lower search costs, they should be able to find lower interest rates, but in this market, even if you control for a lot of observables, those who are searching more seem to be getting higher rates.
And you might say, well, you know, why is that? Well, you know, it's not difficult. The answer is because, you know, there's screening, and screening is informative. Lenders screen, and even conditional on the observables, it seems like they are -- you know, that they are rejecting people differentially, you know, given our data sets, and the people who are rejected more are searching longer, and it's -- and they also have higher reservation rates, because they know they're going to be rejected with higher probability, which makes them, in equilibrium, to settle with higher interest rates.
So I guess maybe I should emphasize here, in the model, it is not always true that you are going to have, you know, search higher -- or you are going to get higher interest rate. It's an equilibrium prediction that seems to be born in the range of parameter values that we have.

And the approval process leads to some sort of endogenous adverse selection at the lender level, and, you know, taking that into account is important because, you know, we don't want to, you know, infer from the fact that somebody's accepting a higher rate, that this person has higher search cost. It could be that this person could be, you know, higher risk credit type as well.
So let me say a bit more about the data. Again, you know, this is some proprietary data that we got through a very resourceful set of co-authors. On conventional loans, it's a detailed, multilevel loan panel. Again, we have these two separate data sets. One is a sample of granted mortgages for which we have very detailed information about, you know, the characteristics and the ex post performance of the loan, so we see delinquency status.
So along with the granted mortgages and their performance, we also have data on the applications and approval status of these loans. Then what we did is we matched these mortgage applicants and the grantees to their credit reports, using data provided by credit bureau, and these credit reports or the credit bureau data have information about the number of inquiries -- the type of inquiries -- credit inquiries that are

257

being made onto these people's records.

And this data set, of course, doesn't have -- has more than just the mortgage loan information. It has all things, like auto loans, student loans, consumer loans, et cetera. There's a lot of things.

So pretty much everybody who has applied for a mortgage here knows the process. You know, there's an application, there's a credit review, and then there's a deposit, and it goes into underwriting at the bank, and then, finally, after 30 or 45 days or, you know, if your seller is somewhat sane, you know, in a relatively short amount of time, you close on the house.

How about the -- the credit review, this is where, you know, a credit pull is done on your report, right? So the bank says to the -- one of the credit bureaus, we are going to do a credit pull, and this is going to be registered as an inquiry.

Now, in this paper, we use a window -- in most specifications we used 45 days, but sometimes, you know, it can be 30 days, all of those inquiries as a proxy for search by the borrower. So you might ask, you know, are all these inquiries mortgage-related? You know, it is possible that, you know, some of these inquiries are done for credit cards or, you know,

258

other types of loans, but actually if you look at the loans that appear in people's credit accounts, this period of about 30 to 45 days before a mortgage is approved is typically very silent, and -- you know, because a lot of people and, you know, their real estate agent or their broker or their friends will tell them, you know, sort of -- you know, you sort of don't want to -- you know, focus more on mortgages here, especially leading up to buying a house, you know, don't do too much searching around. After you get approved, you are going to get a lot more search behavior and other -- for other types of loans in the record, okay?

So once again, okay, once you control for a lot of these covariates -- and my co-authors always tell me, sort of just, you know, say that we have a lot more covariates that -- you know, they argue that some other people have studied the data have used, so -- and we have quite a few covariates here, and we still have, you know, a lot of residual dispersion, and, you know, it stays on even if you control for things like lender fixed effects. So these are sort of very sort of fine-level cuts of the data. You still get -- you are going to get price dispersion.

What about the search angle? This is where we

259

have less other kinds of information. So for the approved sample, the median person seems to search about two lenders, and, you know, below the median, there is only one lender, but there is a tail of people who seem to, you know, search three, four, five, you know, lenders before getting approved. The applicant pull, this is where this very long tail appears, you know, some of these people are -- have huge numbers of credit inquiries on their reports, and, you know, maybe not surprisingly they don't seem to get approved. They don't show up in the approval data.

And the search patterns do seem to certify the creditworthiness of these borrowers. You go from sort of people who are -- you know, have low FICO scores, searching quite a bit more, to people with high FICO scores in detectable ways.

That said, beyond creditworthiness, other types of demographics tend to not come in as clearly. For example, you know, if you do the breakdown by education, you get much smaller differences in search behavior, and, you know, we ran a whole bunch of regressions, and, you know, the number of -- the signs in which -- beyond what the FICO score predicts, how these covariates enter into search behavior doesn't

260

seem to be that interesting or very intuitive in many ways, okay?

So beyond, again, this FICO score difference, it does seem like the evidence on the search and characteristics is a bit mixed and difficult to interpret, so -- but let me now try to put these things in a little bit of model framework, and I will show you the main empirical findings here.

Once again, sort of, you know, the intuition of basic search models, we expect higher search, it tends to be correlated with, you know, cheaper mortgages, finding sort of lower rates, and you might expect some of these characteristics to be correlative with, you know, the sophistication or the search costs of these consumers.

And just to formalize it -- and then I am going to put the equations up, because I am going to modify them for our preferred model -- in the basic model, the sequential search model, there's some search costs that can be heterogenous across consumers, and people get some utility from the mortgage and get disutility from payment, the R-sub-Js, and the lenders are posting rates, and they're competing on rates.

And the consumers are going to follow some reservation rate strategy. If they find -- if they

261

get an interest rate draw that's below their reservation rate, they're going to stop and get that mortgage, and that defines cutoffs in the search cost distribution, which in turn defines the market shares, you know, of the different lenders depending on the interest rates.

So if you simulate data from a model like this, where simply costs are given some bell-shaped, you know, truncated normal like this, and you generate the relationship between interest rates and inquiries, you are going to get a downward sloping pattern where, you know, the rate that you get, it declines with the number of inquiries that you make, precisely because, you know, people who have a higher number of inquiries have the lower search costs and have basically, you know, lower thresholds. They will not, you know, stop until they get the lower interest rate.

So that was -- the previous one was theory, but this is data without controlling for anything. So this is for the approved sample. So interest rates as a function of inquiry show this -- you know, there's a decline apparent in the very beginning, but the U sort of turns in the wrong way when they go to higher inquiry levels, okay?

And this is not just because of -- you know,

262

so -- and this is actually -- you can see sort of how this could be generated due to the conflation of, if you will, different credit types and the search costs, because if you look at the people who have low FICO scores, the relationship is increasing, right? The number of inquiries, higher inquiries corresponds to higher interest rates.

But even for people with relatively good credit, FICO scores above 720 -- and I don't know the population of distribution of FICO scores in this audience, it's probably pretty high, you know, I'm guessing a median of 800 or more or something, so it's going to -- it still shows the U-shaped pattern. It doesn't go along with theory, especially at the upper tail of the inquiry distribution.

If we look at income, even for people who are relatively wealthy or have higher income, assume we have the same patterns, and, you know, across demographic groups, et cetera, this seems to go through, and let me just, you know, again residualize things, not to -- but if you put in a whole bunch of borrower controls, you still get this pattern, you know, once you control for a lot of things, including the FICO score, you get this pattern.

People who search more, you know, are getting

263

higher interest rates, which actually is -- you know, since we're putting in the FICO score as well, it does seem like, you know, there is something being screened on beyond FICO.

So, once again, we want to say credit products are different, and then let me just sort of show you the model, did a simple tweak to the basic sequential search model that's going to generate hopefully the patterns that we see in the data. We are going to introduce some difference in credit quality, and we are going to introduce the screening process that these lenders can reject applicants.

So once again we have a continuum of search cost distribution, but we have a difference in creditworthiness by the applicants. They are, you know, different, and there are two types in this market. You know, Gaston had a lot more types than we did in his paper. We are going to be much less ambitious and have only two types, a high type and a low type.

And this is -- and all of our analysis is going to be conditioned on all these covariates, so this is the residual, if you will, unobservable heterogeneity that affects payment ability, you know, conditioned on all the observables.

264

And the utility for mortgage, we make it, of course, a function of your payment ability, and we can allow for adverse or advantageous selection based on the sign of the Sigma, as in many models. So if the low types have higher utility from the mortgage, you might expect some sort of adverse selection, but in terms of -- in this model, it's interesting, and I don't know how generic this prediction is in other models, that the sign of the Sigma just doesn't matter, you know, in the search model, because everything's based on the differential gains from the next search, the Sigma component just washes away, and what we're left with is a modified reservation rate equation, which is the middle equation here, where, you know, it used to be that I'm equating the search from -- the next search to the expected benefit of the search, and now I just scale the expected benefit of the search with your probability of being approved for the loan. That's the only difference in the model, okay?

But this is what's going to happen if this Pk, which is your approval probability, is low, then that's going to increase your reservation rate, if you will, the R-upper-bar that you are going to try to search for, and so you are going to be willing to

265

accept higher prices, and, you know, in the same amount of time, you are going to search less, and you are going to be -- but because you're approved with lower probability, you may have to search longer. So you have this tension between you having, you know, let's say a worse threshold, higher interest rate threshold, but having to search longer because you are not going to be approved easily.

How the supply side is going to be in this model, and, you know, we have the supply side because we would like to make some counterfactual simulations, and, you know, see what -- how the equilibrium changes in the market, and this turned out to be pretty difficult problem actually.

And, you know, I should talk to David and his co-authors, you know, they are masters in computing these models, you know, carefully, and I love their work because of that. It turns out sort of, you know -- and I'll be super honest about this, I had never written a paper with adverse selection in it this way, and, you know, as the warnings from your first year micro classes might say, the warnings from the first year micro classes were true. These are difficult models.

So it turns out, again, equilibrium existence

266

is not always very easy, you know, markets unravel, you know, there's no -- so what we did, you know, is, you know, this is in some ways a dirty trick, is we put a noise term in the profit function, the size, basically for -- and we discretized the rates that banks can post, which is the empirical reality in this market.

They all seem to be, you know, clustered around one-eighth of a percentage point, you know, offers, and I have no idea why that is. It's a bit like the SEC's stock ticker type stuff, and there might be some interesting, you know, anticompetitive things to study there. But assuming this discrete strategy space, you put this noise in there, it turns into a Bayesian type game where everything's a probability. You know, you can search for a fixed point as if it's a mixed strategy game, and that's how we sort of tried to solve this problem on the supply side.

But I don't want to say, you know, that, you know, that's the end of it, sort of -- it's a tough problem to solve for, you know, supply side when there's adverse selection that goes on in these markets, and it's a very interesting set of economics that goes into it.

So what about the lender? The lender's payoff

267

is dependent, of course, on the rate that they get but also the repayment probability, and the lender makes a forecast of the repayment probability. So basically they have this screening technology, basically this probability is that they get the signal, and the signal is whether there's a high type or a low type, and depending on the signal, they approve or don't approve the loan, okay?

So we simulated some of the data from this model with some assumptions. For example, if you assume that the -- Lambda is the proportion of high types in the population, and about 70 percent good types and 30 percent nonrepayment types, and the -- and we assume that the lenders have very good discrimination ability, so they get basically the high types, 95 percent probable to write, only 5 percent they make a mistake.

So in that model, with search costs being the same across the high types and the low types, what you're going to get is the high-type consumers. The only difference is they're -- you know, the type of being creditworthiness are going to have much lower reservation interest rates, and the low types are going to have higher reservation interest rates, and this is going to yield this upper sloping pattern that

268

I showed in the data that, you know, in equilibrium, in equilibrium rates that are being sent by the lenders, that the lower types, if you will, search longer even though they're willing to settle for a higher interest rate. So the people who are sort of doing a lot more inquiries are getting worse rates or higher rates in this market.

And, again, you know, this is the distribution of types across the different interest rates that are being posted -- given by the lenders. As you increase the interest rate, essentially after a point, after about 2.5 percent, pretty much all of the people you are getting are the low types. So this is really sort of the adverse selection problem, you know, hitting these lenders, okay?

So this simple model generates this dispersion in prices, this positive empirical relationship between search intensity and the prices of the rates, and, again, the only sort of, you know, new ingredient that we put in is this difference in creditworthiness and some sort of, you know, somewhat effective screening technology.

And, you know, we can have other predictions from this model with creditworthiness and screening. For example, we can ask, you know, how do defaults correlate with search intensity and are approvals correlated with search intensity? So in the model, for example, you know, we can generate an upward sloping relationship between inquiries and default rates.

You know, as we might intuit, people who do more search are the lower types, so they default at higher rates, and this is the data that, you know, people who do more inquiries default much more often, you know, even if you control for all these covariates.

How about approval? You know, in our people who are doing more inquiries, are they approved less often? Well, in some ways, they have to. That's what the model says, and that is what the data says, that people who do -- you know, as I said earlier, the people doing a lot of inquiries are approved a lot -- less frequently in the data.

So we might say, you know, so once -- you might also say that, you know, why do you need search in this model? Maybe it's this screening is just what's creating this price dispersion, because, you know, okay, maybe I convince you that there is some extra sort of unobservables in this process that the lenders are able to extract from looking at these applications

and price accordingly. So maybe all of the price dispersion is generated by that.

So to rule that out, what we looked at is the set of what we call never-rejected borrowers. So basically these are people who are very high -- credit scores are very good credits -- credit risk. So these are people, you know, whose FICO scores are basically like people in this audience, above 800, low loan-to-value ratios, that the income ratio is low, you know, very vanilla contract, a 30-year, fixed-rate mortgage.

And for these people, the mean approval rate is about 99 percent, and the -- you know, the relationship between inquiries and approval -- the rates that you get is upward sloping as the standard model predicts, and the -- sorry, this should be the other one. The number of -- it should be downward sloping as we predict, so this is the -- I'm trying to think what's the left one, but -- so this is all borrowers, so that's all par borrowers is upward sloping.

For the never-rejected sample, it's downward sloping, that the people who sort of are searching more are getting lower interest rates in this never-rejected or very good credit sample. And, you

know, you can do it other ways. You can do it by logit score, you know, look at the default rate predicted and, you know, give the logit score to these people and find the people who are above 97.5 logit score, and for these people as well, then the interest rate you get is, you know, at least on most parts of the curve declining, maybe weakly declining in the number of inquiries.

So what we would like to say from this is that search does matter for these people, and, indeed, even for this sample, there is quite a bit of search going on. There's a lot of dispersion in the amount of search that people are conducting. Some people tend to do a lot more inquiries than others, and they get, you know, better rates. So what I learned from this is I should, you know, ask more banks for quotes, and maybe I'll get a better rate down the line, but the -- yeah.

So to wrap this up, you know, we have this model, again, that explains this nonmonotonic relationship between, you know, search and rates, and then, you know, I hope I was able to convince you that search is somewhat important, but also the screening and the importance of unobservable risk types or adverse selections are also important in this market.

So now that we have some facts to motivate the model further, what we want to do is to sort of estimate some model parameters maybe to get a handle on how effective screening seems to be in this data, you know, as it fits the moments that we observe in the data, and maybe use this model to do some counterfactuals.

So, once again, sort of, you know, one -- there's quite a few papers in the literature, and I'm guilty of a few of them. You know, what we try to do is we look at observed price dispersion and distributions or both prices and quantities and try to infer demand parameters, which in this case are search costs, but the issue with this in this market is, you know, what all those papers and techniques will give you is the left-hand side of this equation, and the right-hand side is what you see in the data, dispersion of rates, et cetera. That's the theory of the first order condition, if you will, of search models, and the left-hand side is the search cost that rationalizes what you see in the data.

But because we don't have the approval probability in the denominator, we are going to get the wrong inference on the search costs that we observe in the data, if approval is an important part of this -- of these markets. So what we have in this data, you know, because we have a setting with unobserved types, unobserved credit types, how are we going to get at that approval probability?

Well, we have approval data, but we also have this mixture of people, you know, high types and low types, you know, we can generalize it to have more types, but we decide to stay with two types of creditors -- of borrowers. So we have the search information. We have the -- also the mortgage's performance down the line, which allows us to get a sense of what type of borrower this is, you know, conditional on getting the mortgage, and to estimate the parameters from the data.

And the parameters are somewhat interesting. So they seem to indicate that, you know, screening is informative, that, you know, the banks are able to get the high type, you know, so -- so -- so I'm trying to -- 80 plus 2X, and so X is 10 percent, so it's about 90 percent probability of getting the high type right, with 10 percent, you know, mistake in getting the high types right. That's what the 79 percent means.

And there seem to be quite a few bad risks in this applicant pool. If you will, about 50 percent of

them are actually sort of -- should not qualify for this. They are, you know, people who are not going to repay with high probability.

There is some default by high types, you know, as the model classifies the borrowers. About 90 percent of them, you know, default in the data, but -- and so 10 percent of them default, but 70 percent of the bad types will default in the data. That's what the model yields.

What about the search costs? Well, the search -- even if you account for this creditworthiness heterogeneity, it is substantial, you know, it's about 27 basis points. If you try to do it by year or over the 30-year life of the loan, it's about $10,000. So you are basically paying $10,000 more over the life of the loan because, you know, you're not searching one more lender. And then there is heterogeneity in search costs, and the percentiles are somewhat different.

And these numbers are broadly consistent with other findings in the literature about search in markets where credit -- financial product markets where approvals are not that important. J.F. can, you know, say otherwise, but I think in their market, this wasn't -- you know, those were sort of -- approvals

were not that important a feature of their market, so they find other -- they found similar order of magnitude search costs, and in other financial product markets I know of, these are similar types of numbers.

And the model, even though it's very stylized, does fit the -- you know, the basics of the data relatively well, and so the -- we are still -- you know, the reason this paper is not out the door is because we haven't done as much as we would like in counterfactuals, but I will show you what we have so far.

So one thing we wanted to do, my colleagues being more sort of finance macrotypes, they were interested in how sort of monetary policy changes are transmitted in the mortgage markets. I said I don't know too much about that, but they wanted to look at ten basis points in reduction of cost, how is it transmitted in this market? Was it passed through?

Essentially the answer is, in this model, it's about one-for-one pass-through, so it's a -- you know, even though there's a lot of sort of, you know, search costs, you know, all this adverse selection goes on, pass-through still seems about one for one.

Another one that's maybe a bit more interesting is a calculation that is counterfactual regarding

redlining practices. So this is a very nice picture that Gregor found in some court records about this bank, Evans Bank in Buffalo. This is a court document. They actually sort of, you know, redlined the areas where they operate, and the hashed lines are the places where the population is very -- majority or near majority are African-American, where they do not operate.

So -- and there was a lot of redlining lawsuits of this kind, and what we tried to do in this simulation is, you know, instead of doing explicit -- so instead of doing discrimination on rates, based on race, we are going to do the discrimination on the approval of these creditors. So basically some of these banks are going to systematically approve the applications of certain kinds of applicants with much lower probability than others, and -- which means basically that approval probability is going to be sort of penalized by this discrimination factor.

So what's going to happen is that the discriminated group in this model realizes this or learns about this once they've done a few applications, so that they are going to approve with less -- with lower probability, so they are going to search longer, but they are also going to raise their reservation interest rate. So they are going to settle for a lower mortgage.

So what this is going to -- what this is going to do is even if there's no explicit discrimination on the rate, because this group is essentially acting like -- you know, more inelastic demand, the rates that these banks are going to offer are going to rise, and because in this model we assume that the consumers do not know which bank is discriminating, and they are going to approach each lender with the same higher reservation interest rate, so what's going to happen is that the nondiscriminating banks as well are going to charge higher loan rates in equilibrium in this model.

So what's going to happen is that the overall interest rates are going to go up quite a bit, the average interest rate, and actually sort of the amount of searches that, you know, that have to be done in the market is going to rise as well. So the mean origination rate is going to -- and it depends on the percentage of people who are redlined against, of course, and then these other parameters, but, you know, I think -- and the interesting aspect is the strategic complementarity, if you will, if the discriminators, but the redliners are, you know,

cutting down on the approvals, and all loan rates go up for a -- against this discriminated group. So essentially the market discriminates against these people through this effect.

Other things we did, you know, what about tighter lending standards? So we definitely see in the data, when we do it by subsamples, that the shift in these screening probabilities -- one anecdote to motivate this is, you know, we read somewhere that Ben Bernanke was rejected for a refinance loan, you know, at the -- near the height of the crisis, so there was a time where -- you know, after 2008 where banks got very, very sort of conservative, if you will, in their screening practices.

It does sort of affect, you know, people's search and acceptance probabilities, reservation rates quite a bit, and it increases the interest rates by about -- by some, and the model search introduced -- that's done by these people, which is definitely the people seem to be searching more during the crisis times in our data.

On the reverse side of it, there are policy interventions, like the Community Reinvestment Act, which is basically regulations that weaken strict screening technologies. These are restrictions on

screening. And you can try to do exercises where you may see the shutoff screening, and, you know, we look at the effects of this type of policy.

An interesting sort of overall message that we get from it is, you know, the changes in equilibrium offered rates especially come from two sources. One is the demand side adjustment, you know, people changing their reservation interest rate, but also how the supply side reacts by the offers that they give. It appears that most of this, you know, when you shut down the supply response, if you will, and you look at only the effect of the demand side, you don't get as much movement in these equilibrium quantities. The supply adjustment component is much more important quantitatively than the demand side effect on these equilibrium outcomes, okay?

So let me stop here. There's a zero there, and so I just want to say again, you know, search has been a very fruitful area, you know -- of course, since 1961, I saw about 9000 Google Scholar cites on Stigler's paper, probably many more, on, you know, explicit citations. You know, a lot of people think about search models.

What we want to do is here, you know, in these credit markets or financial products markets,

especially the screening or approval process has also become an important aspect of it, and we need to take this into account.

Again, I would like to push this -- you know, I am going to try to do it, but if people are interested, I think the -- how these institutions do screening is very important and how sort of the regulations affect these screening technologies, how, you know, these organizations use data to screen people is also very important, you know, in determining equilibrium outcomes.

So, you know -- and then, again, sort of not just credit markets, but lots of insurance markets, you know, fall under this point of view, and sort of I think, you know, over the last decade or so, there's been a lot of work on markets with -- you know, markets we call selection markets or markets with adverse selection, bringing the theory and empirics together, but, again, I want to say there can be nontrivial implications of this in the data and also in the -- you know, in the execution of these, and I think it's going to be, you know, very interesting and challenging times for, you know, applied economists for some time to come.

Thank you so much for inviting me.

(Applause.)

MR. ROSENBAUM: Thank you, Ali.

We have a reception outside, and we will continue the conversation there. Thank you all very much.

(Whereupon, at 5:20 p.m., the proceedings were adjourned.)

CERTIFICATE OF REPORTER

I, Jen Razzino, do hereby certify that the foregoing proceedings were digitally recorded by me 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; that I am not a relative or employee of any attorney or counsel employed by the parties hereto, not financially or otherwise interested in the outcome in the action.

JEN RAZZINO, CER
Court Reporter

E

G

In the Matter of:

The Eleventh Annual FTC Microeconomics Conference

November 2, 2018
First Version

Condensed Transcript with Word Index

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

UNITED STATES FEDERAL TRADE COMMISSION

THE ELEVENTH ANNUAL
FEDERAL TRADE COMMISSION MICROECONOMICS CONFERENCE

Federal Trade Commission
FTC Constitution Center
400-7th Street, S.W.
Washington, D.C.

Friday, November 2, 2018
9:00 a.m.

Sponsored by:
Federal Trade Commission Bureau of Economics
and
Northwestern University's Searle Center on Law,
Regulation and Economic Growth

C O N T E N T S

P R O C E E D I N G S

\- - - - -

WELCOMING REMARKS

MS. DUTTA: Hi, everyone, and welcome to the first session of the second and final day of the FTC Microeconomics Conference. This session was organized by Katja Seim of the University of Pennsylvania, who's a member of the Scientific Committee for the conference this year.

As you may have seen with the sessions yesterday, there will be two papers presented during the session. For each paper, the presenter will have 25 minutes to present the paper, which will then be followed by the paper discussant, who will have 10 minutes, and then finally we will have about 10 minutes for Q&A.

(End of Welcoming Remarks.)

PAPER SESSION:
HOW ACQUISITIONS AFFECT FIRM BEHAVIOR AND PERFORMANCE:
EVIDENCE FROM THE DIALYSIS INDUSTRY

MS. DUTTA: So I would like to invite Ryan McDevitt of Duke University to present the first paper of this morning, which considers the effect of acquisitions on firm behavior and performance in the dialysis industry.

Ryan, welcome.

MR. McDEVITT: Well, thank you very much for letting me be on the program. It's very exciting to be here. It's a great conference, great audience, although this reminds me of an MBA classroom at 9:00 a.m. There will be some stragglers coming in, I suppose.

This is joint work with Paul Eliason. He just started at BYU as an assistant professor. He's doing great work, a lot of it in the dialysis industry; and Ben Heebsh, he's a current Ph.D. student, and he will be on the market next year, I believe somewhere in health economics. He will be a great person to look at. And then Jimmy Roberts is also on this paper.

The main motivation for our paper today is we're going to look at the consolidation in the healthcare industry. As we all know, there has been rampant consolidation over the past few decades, and as IO economists, we're in a great position to analyze the effects of this consolidation.

The current state of literature in IO mostly focuses on the effect of concentration on outcomes like prices and quality. Typically more concentrated markets have higher prices and lower quality, but the literature that we're aware of, it mostly looks -- its concentration is somewhat of a black box and how that affects the outcome. There's some measure of HHI on the right-hand side, and then these outcomes are on the left-hand side.

Our talk today will focus more on what's going on behind the scenes. How does the firm actually move from an acquisition to effecting outcomes? We want to dig in, in a very precise way, and understand how these prices and how this quality changes after a merger/acquisition.

And we came upon this topic based on two of our previous papers. In long-term care hospitals, I worked on this setting with some of the co-authors on this paper, and long-term care hospitals specialize in patients who have very prolonged needs. They have been in a car accident and they need assistive breathing; they are in the hospital for several

months. And Medicare has a quirky reimbursement system where they give a short per-day reimbursement for the first few days of the stay and then a lump sum that's supposed to cover the whole length of the stay after about two or three weeks.

Based on this compensation scheme, long-term hospitals tend to discharge patients right after they reach that lump sum threshold, which obviously distorts care that's not efficient from the patient's perspective, but we were very intrigued by the result that after the long-term care hospital chains, there are two big chains in this industry. Once they acquire independent hospitals, they tend to implement this strategy more often. It seems like there's some kind of transference of best practices from a profit-focused standpoint.

We then can combine that with a paper I worked on with Paul Grieco, where we look at what we call the quality/quantity tradeoff in dialysis, where some for-profit chains tend to emphasize quantity over quality, because it's more profitable to do so, and so synthesizing these two results will be the main focus of our paper today.

To do our analysis, we're going to look at 1200 acquisitions that have taken place over the past 15

years in the U.S. dialysis industry, and this is a great setting for our purposes because the large chains here, DaVita/Fresenius, they behave very differently than the independent facilities.

They use more injectable drugs, for instance, because they're very profitable during the time period of our study; they replace nurses with techs because nurses are more expensive than techs; they treat more patients per employee and station, trying to be more efficient, stretch their resources.

And in doing this it leads to worse outcomes for patients. We find that survival and transplant rates fall once an independent facility is acquired, and hospitalizations increase. This, of course, wastes Medicare's scarce resources. Medicare is paying more for lower quality outcomes.

There has been much work on this topic, both within healthcare and outside, by IO economists. I can't spend any time on this really, I have only 25 minutes, and this is not an exhaustive list by any stretch, but we think of three main buckets of literature right now on this topic.

The first is that looking within healthcare and even in other industries, typically you don't consider mechanisms, but how quality and prices change. Again,

they look at a very reduced-form way of how concentration then affects the outcomes. So we're going to build on that by looking precisely at the mechanisms.

There has been some work on how firms transfer their strategies after an acquisition, and there's going to be evidence here that managers implement these best practices within dialysis. And, of course, within healthcare, the reimbursement scheme for Medicare matters a lot for providers' behavior, and we're going to show, just like we did with long-term care hospitals, that these providers respond to incentives. Not a mind-blowing result, but we can show precisely how this happens.

I'll give you a few brief details about dialysis. It's an important industry within healthcare, covers a lot of patients. A patient who needs dialysis because they have kidney failure -- and the kidneys perform two main functions within the body. They filter toxins from the bloodstream and they also stimulate the production of red blood cells. If you have kidney failure, the kidneys no longer perform those functions, so you need a replacement to continue to live.

The two options for this are dialysis, where you see in the pictures you go into a facility, you're hooked up to a machine, and that machine replaces the function of the kidneys, it filters blood and toxins, or you can receive a transplant. That's the most preferred option. That's the only way to actually cure this condition.

The issue, though, is that kidneys are scarce, there aren't enough to go around, and so practically speaking, all patients with kidney failure at some point go on dialysis.

In the United States, dialysis is an outsized influence on spending on healthcare. There are about 500,000 patients across the United States, and 90 percent of these are covered by Medicare. In the seventies, Congress enacted legislation that covered kidney care in the United States for all patients regardless of age. In John Oliver's segment on the dialysis industry, he made the joke that it's like one organ of the body in the U.S. is Canadian, the kidneys, because we have universal coverage.

There's an 80/20 split with Medicare Part B, so patients pick up 20 percent of the costs, and if they have private insurance, that covers the first 30 months. And this will be an important feature of this industry, because privately insured reimbursements are

much larger than Medicare reimbursements.

But the bottom line is, we spent over $30 billion a year on this, at 6 percent of Medicare's budget and actually 1 percent of the overall federal budget. This is a huge issue and it's growing considerably over time.

Now, I'll briefly go through Medicare's payment structure for dialysis. During the time period of our study, centers paid a composite rate of -- were paid a composite rate of $128 per treatment, up to three times per week, but drugs like EPO and some other injectable drugs were paid on a fee-for-service basis.

Of course, this led to some wasted resources. Centers put too much EPO into patients, so Medicare reformed payment in 2011, and now it's $230 for treatment and drugs within one bundle. And you can see from this figure that after the bundle reform in 2011, use of injectable drugs fell considerably, which to us suggests that the behavior of centers is really influenced by this payment structure.

EPO is going to be a main focus of our paper. It treats anemia. It's used by 90 percent of dialysis patients at any given time, and it was the largest drug expenditure for Medicare for many, many years, almost $2 million in spending in the late 2000s.

Facilities would get $10 per 1000 units in reimbursement, and this added up to 25 percent of DaVita, one of the largest chains in dialysis, of their revenue, and 40 percent of their profits. So this is a huge profit center during the time period of our study.

The structure of this industry, there are about 7000 facilities across the United States, and growing more each year, and two large chains that dominate this. So think of this as a duopoly, DaVita/ Fresenius, two for-profit chains, and despite their claims in the press that they aren't reimbursed enough to actually cover their costs, they're very profitable, and those profits have been going up over time.

And to give you a sense of how this industry has evolved over the past decade or so, we can see the growth in facilities but also that DaVita and Fresenius are becoming more concentrated. They own now up to two-thirds of all facilities, and a lot of that has come through acquisition.

And here's the plot of acquisitions over time and how they've grown. The bottom dark blue segment of this figure is independent acquisitions. That's what we're going to focus on in our work. The big

spikes come from large acquisitions. We don't consider those in our analysis today, because we think there's other issues going on when you can acquire a big chain, trying to integrate that big chain. It's the independent facilities that we can really focus on the details of how they transfer firm strategy.

And we do have quite a large sample of acquisitions, which is one of the reasons we're so intrigued by this setting. We're going to use over 1200 over time, and there's a lot of variation over years. And this is going to allow us to identify some of these effects. Because we have such a large sample of independent facilities and great data, we can then use this to analyze what's going on and then show you at the chain level how these acquisitions are occurring, right? DaVita and Fresenius have grown considerably through acquisition.

And another point to make is that strategy seems to be very important for dialysis chains. Kent Thiry is the CEO of DaVita. He's a well-known figure in healthcare. He is a Harvard MBA. He's known to be a little bit eccentric. He thinks that culture really matters for DaVita. His motto is in healthcare, it's so important that it's all for one, one for all, and he dresses up like a Three Musketeer for some of their company events to really make that point.

And so I'm just trying to demonstrate in this slide that strategy is very important, culture is very important for these facilities, so that's a channel through which facilities might change their behavior after acquisition.

So our measures for the effects of acquisitions, we're first going to look at observable provider choices, aspects like injectable drugs, EPO, for instance; we'll get staffing decisions, whether they have nurses or techs; we'll look at the overall staffing level, how many resources they put into the facilities. We'll then see how these influence clinical measures like what we call the urea reduction ratio, how much of their toxins are cleaned through dialysis; and also hemoglobin, what's your blood level like after you get injections of EPO.

And then we'll also look at patient outcomes, factors like hospitalizations, mortality transplants, and that will allow us to also measure some aspect of quality.

And the reason we can do any of this with our paper is that we have really incredible data for the dialysis industry. Because Medicare is the primary payer for all dialysis patients, they make all the

data available to researchers, so we have over 14 million patient months at a very detailed level. Every month a facility must file claims, and we have the claims data -- nonitemized, of course -- but we see for each patient, for instance, how much drugs they receive, what kind of treatment they receive, their blood measures, their infection rate, their hospitalization rate. Everything that we would want to measure, we have access to that in the data.

To give you a sense of what's in our data, here are just some selected summary statistics broken down into four categories. We think of facilities as being always independent, and then the independent acquired facilities, we look at them before and after acquisition, and then we also have facilities that are always a part of a chain. And you can see from this table, there are really noticeable differences, at least in an observable way, across these four categories.

And some that pop out are really due to the time series, just of evolution and trends over the time period. For instance, ischemic heart disease has fallen considerably across the U.S., and clearly because we have a post-acquisition dummy, that sample is from later periods when heart disease has fallen.

And so really, instead of a summary table, we're going to show you the regressions, and those will be much more telling of what's going on.

Identification strategy is very straightforward. Think of the simple dif-in-dif, where we're going to look at how an acquisition affects outcomes, and the two primary threats to identification here will be first it could be that patient mix changes after acquisition, and so it's not the acquisition itself that changes outcomes, it's just you're looking at different types of patients, and for that we're going to rely on very robust clinical and patient data to understand how these effects are changing.

And the other key issue is that obviously acquisition is not random. These chains are picking off facilities, and to control for that, we're going to include facility fixed effects, which will be crucial, because we're looking at, within a facility, how behavior changes after acquisition. That means identification is truly from within physical changes in ownership, and we'll also show you there's no trend prior to acquisition. So we're okay in a dif-in-dif sense.

And our advantages here in this setting over

previous studies, first we have a very large sample of acquisitions. 1200 is the largest we've seen. Of course, we would be happy to see other papers that also work on this, if we haven't covered them yet, but 1200 is a very large number for acquisitions.

We also have cleared channels through which strategies could change after acquisition. There's a limited scope or change in prices here because Medicare unilaterally dictates reimbursements. There's not much going on in terms of price competition. And there's little evidence here that market power matters, at least for Medicare patients. We'll show you some results at the very end, and that's more the work of Paul Eliason, but it really is to worry about firm strategy, not about competition.

And here is the main figure for the paper. If you gave me only one slide to present today, this is the slide I would show. In this figure, we have EPO dosing at acquired firms, in the left-hand side of the panel is months prior to acquisition, right-hand side is months after acquisition. And you can see clearly there's no trend before acquisition, very flat, this is normalized coefficients. It's very flat EPO dosing.

And then after acquisition, a very sharp

17

increase in EPO doses, which can't be explained by clinical necessity. It's purely the result of firms seeking profits. For us this was a very stark finding and really what we're going to build the rest of the analysis around.

That figure comes from a regression, and this regression, we can think of it in a few different ways, but I think the most restrictive, most conservative regression is in column 4, where we have a host of controls plus key fixed effects, including year/month fixed effect, patient facility controls, facility fixed effect, in addition to patient fixed effects.

So that means identification is coming from within a patient, after a facility is acquired, how does that patient, himself or herself, change in terms of EPO? And so that's a very conservative regression and we're very confident in these results.

Another injectable drug to look at, Venofer and Ferrlecit, these are iron supplements. People on dialysis are often deficient in terms of iron, so they receive an injectable drug. And here you see a clear pattern where, after acquisition, the use of Ferrlecit drops and the use of Venofer increases. And the explanation here is that Venofer is reimbursed at a

18

higher rate, even though they're perfect substitutes, just some quirk in the packaging and the size of the vials they use. And so, again, it's a clear profit motive. If you use more Venofer, your profits will go up, even though from the patient standpoint, they're equivalent.

In terms of resources, we look at certain ratios, for instance, nurses over techs. Nurses are higher skilled, but they're more highly paid; they have higher wages. And we see that the nurse to tech ratio is about one to one before acquisition, right? After acquisition, it falls 15 percent. So it appears as though the for-profit chains' nurses -- they substituted techs for nurses because it cuts their costs, and potentially that will have an effect on outcomes that I will show you in a moment.

They also stretch the employees by putting more patients per employee. The patient-per-employee ratio increases by 12 percent. And they also have more patients per station. Patients per station goes up 4 1/2 percent, and it's going to be very bad for dialysis, because patients per station, for instance, means that they have more turnover on each station, which means they have less time to clean the machines between use.

19

And because these patients are hooked up, their blood is exposed to a machine, that means they're susceptible to infections. If we don't clean it thoroughly, that means a higher turnover makes it a greater risk for infection, and this is borne out in the data. We find that patients at acquired facilities mostly fare worse after acquisition.

For instance, all cause hospitalizations go up 6 percent, and again, this is very (indiscernible). Looking at this very same patient, before and after acquisition, what happens. Their risk of going into the hospital goes up 6 percent. Risk of a blood infection goes up almost 3 percent. This is one of the most severe conditions you can have, very hard to recover from, very painful, very costly to Medicare in terms of hospitalizations, and, again, the story here is that because they have more patients on each station and fewer nurses and techs to clean the machines, they're at greater risk of acquiring a blood infection.

Also, EPO doses at too large a dose increases patients' risk for a cardiac event, and we see those go up almost 4 percent. And, again, this is a very bad outcome for patients. They're at risk for this, and we see because they're getting doses of EPO that

20

are too high, their risk of a heart attack goes up.

We can also look at less acute measures from clinical outcomes from the dialysis itself. Good URI is probably the one measure we find where there's unambiguous increases in quality after acquisition. Patients with good URI, meaning their blood has been cleaned of more toxins, that goes up 2 1/2 percent after acquisition.

Low hemoglobin falls because of all the EPO, that's a very small change, even though statistically significant. But on the other side of that, high hemoglobin goes up by 4 percent, which is bad in the sense that it increases the risk of cardiac events. And good hemoglobin within the recommended range, that falls by 3 percent.

And probably the most important statistics for patients is how likely they are to survive dialysis or get a transplant, and based on our analysis and both measures, patients do worse after acquisition; less likely to be on the wait list to receive a transplant within the first year, that falls 9.4 percent.

And again, a transplant is the only way to cure this condition. It's the most preferred outcome, most preferred treatment option for kidney failure, but a tradeoff for a facility is if someone gets a transplant, then they're no longer a customer for the dialysis facility, so they have a conflict of interest there.

And there are some lawsuits that look at just that issue, where DaVita/Fresenius have been accused of not promoting transplants or being on the wait list for their patients, which conflicts with federal guidelines.

Patients are also 1.7 percent less likely to survive their first year of dialysis. Mortality rates are higher after acquisition. Again, a very bad result for patients, I think it should go without saying.

And then the bottom line number for Medicare, payments go up about 7 1/2 percent after acquisition, and this is what the facilities are trying to implement with their strategies. They're profit-maximizing entities. They want reimbursements to go up, and they've achieved this mostly through drug use, but on the cost side as well, we see the costs decline after acquisition. So revenue up, costs down, profits are going up considerably at these facilities.

So to conclude briefly on -- I can spend some time on this slide, but the bottom line from our study is that acquisitions lead to worse outcomes for

patients, higher reimbursements for Medicare, which means the overall value of these treatments have unambiguously fallen, where the payers are paying more for worse quality of care, a very poor result.

And one aspect of our study that I didn't spend much time on today is that there's not much evidence that competition matters in dialysis. We think of these facilities as being their own individual local monopolies, and Paul Eliason has spoken on this extensively in his job market paper, because these patients are in very poor condition, often very low income. They have very high travel costs to get to a facility. So there's very little switching that goes on regardless of quality.

So once the quality falls at the acquired facility, there's not much response from consumers, which is puzzling if you think of free choice here and you are free to choose any facility that's available, but they don't switch because travel costs are so important. They're almost always going to one that's closest to them. We see fewer than 1 percent of patients switch each year even when quality falls dramatically.

So the next part of our study, we're going to really focus on this competitive aspect and try to

understand why competition doesn't matter in this industry. Our hypothesis is it's because of these travel costs, but it's something we want to spend more time on as we revise the paper.

We also have a number of future projects we want to work on in this industry. The first is a study of EPO use after bundle reform in 2011. I showed you the figure where EPO use fell considerably right after payment reform, which, again, is not surprising, because reforming the bundle meant that EPO went from pure profit, they got a markup over the wholesale cost of the EPO drug, but after bundle reform, that became pure cost because it was part of the bundle. So pure marginal cost, which means as the firm is trying to maximize profits, they will use less EPO.

We are going to look at that specifically in another paper, and here we have a great potential instrument. The elevation of the patient affects the size of their EPO dose. At higher elevations, your blood just naturally produces enough red blood cells, you naturally have enough red blood cells, and so we use that instrument to understand who will be more affected by a change in payments for EPO.

The second paper we want to write on this

setting looks at what we call the make or buy decision for facilities. These chains have acquired a number of facilities that we see in this figure, but they also do a lot of new investment, and there is even a little bit of exit. And so we want to understand how access is affected by payment reforms.

One counterpoint to all this is that maybe we wouldn't have any facilities at all if they weren't allowed to earn such profits from cutting quality; we don't see much evidence of that. And another argument against that is that the U.S. outcomes are much worse than in other industrialized nations, which shows it is possible to have this industry without such payment reforms, but we want to focus specifically on a more structural model to understand when facilities enter a market and how that's influenced by payments.

Thank you very much, and I'm looking forward to the discussion.

(Applause.)

MS. DUTTA: All right. Thank you, Ryan. The discussant for this paper is the FTC's very own Nathan Wilson.

Nathan?

MR. WILSON: Well, thank you very much for your attention, and thanks, everyone, for coming out. I want to start, as the other participants have, by thanking the folks who put me on the agenda this year. I was not one of them. I just said yes. So no nepotism necessarily involved here.

Now, before I get to talking about Ryan's excellent paper, I have to start with the standard disclaimer, that the following views are solely those of myself and they do not necessarily represent the Commission as a whole or any of its constituent commissioners.

Now, I want to start by just kind of doubling down on some of the stuff that Ryan talked about, which is this industry is exploding, right? Between 1980 and 2010, we had a roughly fivefold expansion in the number of treatments being provided to Americans. Just enormous. And if you look at the U.S. RDS data which tracks, you know, the facilities providing these services, you can see that since around 1990, right, like where we had roughly an equal split between for-profit and nonprofit facilities, that's really diverged, and pretty much all of the growth has been in terms of for-profits.

Well, you know what could explain that, maybe they have much lower costs of capital, it's much easier for for-profits to come in, or maybe their

marginal profits per period are just way, way higher, either maybe for some sort of socially benevolent reason, lower costs, or maybe because they're -- maybe they're -- maybe their lower costs are reflective of lower quality, or maybe they're kind of gaming that compensation system and really squeezing the -- kind of the per unit compensation. That's certainly a plausible story that could explain these patterns.

I want to put up another graph that Ryan showed, which is just kind of the pattern of acquisitions that's been happening over time, right? So we don't just see this changing market structure due to differential entry, right? We see actual acquisitions by existing players of other existing players. So as IO folks, we think, well, yeah, obviously you can't ignore other stuff that could be changing around these deals, but gosh, if we had really great data, we could look within these facilities to see what might be happening in them, and maybe that will give us some insight into whether the changing market structure reflects lower entry costs or alternative behavior, kind of within each kind of period, that could kind of motivate greater for-profit activity.

And so just to kind of quickly resummarize

Ryan's conclusions from the paper, you know, they're doing a very straightforward analysis of like what happens when a for-profit chain acquires an independent, you know, in terms of the strategy they pursue at that facility. It leverages just preposterously sort of pipe dream, fantastical data in order to do this.

I've played a little bit with related data myself. They are great. So because of the high presence of Medicare, you really see just about all of the patients. Because Medicare is the overwhelming payer, they're tracking stuff at the facility level in a granular way that is extremely rare to encounter. So just a lot of fun from a pure, you know, empiricist's perspective, and the richness of those data really allow us, I think, to be very confident in the plausibly causal nature of the effects that Ryan and his co-authors are finding.

And it's also just a model paper in terms of performing very straightforward econometrics to get at elements of interest. And the evidence, as he was describing, you know, shows that patient health, pretty consistently across a wide measure of different outcomes, declines following these deals. And I think really nicely we're able to see, by looking at what's

going on on the clinical side, what could be explaining that deterioration.

We can see that there appears to be shirking on kind of quality inputs, and we perhaps maybe think that although staying within recommended clinical guidelines, maybe the excessive usage of EPO may be associated with some of these negative health outcomes, too, right?

And I think it's always nice to be able to compare what we're seeing in the data to, you know, qualitative stories. You know, just obviously we should trust the systematic results, but it's nice to be able to tell a story. And so if you just do a kind of a cursory Google news search associated with fines and lawsuits associated with some of the major players here, well, you see a lot of stuff that says, oh, these results really pass the smell test.

So, for example, restricting myself to just one firm, in less than an afternoon's worth of Googling -- or alternative search engines, no endorsement being offered here -- I found, oh, this firm paid almost $400 million in terms of improper kickbacks; paid almost half a billion dollars for excessive usage of Venofer; paid $55 million for excessive use of EPO; almost another half billion for Zemplar; almost 100 million for submitting false claims.

At that point, I thought that's probably enough; I've made my point. There's plenty of reasons to think that there is worrisome quality investment here by these chains.

So obviously I hope no one was kind of holding their breath about my opinion on this paper. It's fantastic. It's nice to see that it's -- that others share that opinion. It's R&R at QJE. I think that makes a ton of sense. It's an important topic, well written, nice usage of data visualization techniques. So fantastic stuff.

You know, because we're economists, there's always things we can kind of point at and pick at and suggest them to spend their time on. These were things that struck me as potentially worthy of additional consideration, either in this paper or perhaps in a future work.

So one thing that struck me is I think it makes a ton of sense to focus on the independent acquisitions. They are cleaner in some sense, but if you look inside the paper at who's acquiring these independents, well, it's actually by kind of the nonbig -- disproportionately by the nonbig two chains, who are themselves going away over time. So that's

kind of an odd thing.

It might be interesting to see, you know, is there heterogeneity in there? If we focus just on acquisitions by the big two, what do we see? What explains why the smaller chains are going away? Maybe they are, you know, less profit-minded than the big two. That would be a super-fascinating thing to observe. I think that there might be something there.

In addition, you know, I think Ryan already alluded to an interesting kind of thinking about the competition stuff. I think -- and I've -- candidly, I've written on this, that I think there is things to think about in terms of local market competition. I think the stuff that the paper does is entirely sensible, but I wondered about, you know, what if you started restricting your attention to, you know, more homogenous kind of market areas, so at least sort of comparing urban areas to urban areas, and then, you know, potentially endogenous measures of competitive intensity to see if those results hold up.

So, you know, there's certainly potentially, you know, no impact of local market competition, but it would be nice to see a little bit more work there at some point, maybe not in this paper, maybe in subsequent versions.

And then, of course, just given all the dynamism of market activity here, just understanding a little bit more about, you know, what else is going on when independents are being acquired. Again, I would not expect that at all to overturn the qualitative results, but it would be interesting to see more on.

And then, again, just kind of a final nitpick, extensive margin effects. You know, I think the evidence is pretty clear that, on balance, there are major things to worry about with some of these acquisitions, but maybe there's a story to be told about kind of growing the overall market with these firms through outreach advertising, through outreach of some other sort. It would just be nice to check this out a little bit more so we can be even more confident in our overall welfare conclusions.

So with my final 30 seconds, my big kind of take-away here is, I think, what is so weird about dialysis? You know, we've looked at a lot of other healthcare industries, you know, we don't see necessarily nonprofits behaving exceptionally benevolently in the case of hospitals in particular. If you think about your high-priced markets, where you have quasi-monopolists operating, guess what? Those are nonprofit systems.

There definitely don't seem to be concerns about the usage of exploitation of market power in the search of profits there. What is so unique about dialysis that, you know, these patterns don't recur or, rather, the patterns in dialysis are so starkly different?

So I don't have any other conclusions, but I would really like to see more work done to try and answer that question. Thank you.

(Applause.)

MS. DUTTA: All right. Well, thanks, Nathan.

We are going to have about ten minutes for Q&A. So I'm going to welcome Ryan back to the stage. Thanks.

MR. McDEVITT: Are you managing the questions? Are you managing them? Okay, great.

MR. BRUESTLE: Hi, Ryan. Stephen Bruestle, Federal Maritime Commission.

You've done a good job of showing patients are worse off and facilities are better off due to the acquisitions. Any idea of whether -- and this gets to one of the comments by your reviewer -- any idea of whether society as a whole is better off? I realize this might be a big win because you might have to rely on estimates of the statistical value of life.

MR. McDEVITT: I think from our perspective, quality has fallen considerably, which is worse for society. You know, if we're agnostic about consumer versus producer surplus, it's a little hard to say, but if we're taking it from the perspective of maximizing well-being, then clearly this is worse, because patients are more likely to die, quality of life is falling, and I think there's no evidence that we're expanding access on that extensive margin that you mentioned.

I don't think there's much evidence that this is allowing more patients to be treated. So I think from an overall societal standpoint, I think this is clearly bad for society.

MR. BRUESTLE: (Off mic.) Well, but I also would consider (inaudible) profit and more money into the economy as a benefit. I mean (inaudible).

MR. McDEVITT: Yeah, there's not much we can say in a very general equilibrium setting of, you know, how does the whole healthcare system benefit, but I think really what we're doing is we're transferring profits from Medicare and taxpayers to for-profit chains who are not making the best use of these resources.

I think if we invested the same amount of money

in other types of care, we would be much better off, but I'm speculating there. We don't have a model for that certainly in the paper.

MR. BASKER: Emek Basker, Census Bureau.

I'm curious about whether you have any data about employee turnover or anything like that. I can imagine if you're working in a facility that's starting to change its practices in ways that you might find very unattractive, that that would be one metric of what's going on.

MR. McDEVITT: Yeah, a great point. From the data we have, we have fantastic data, but we don't have, at an employee level, who the actual employees are. We have measures of, like, how many actual nurses and techs are employed at the center, but we don't see a turnover measure.

But I showed you that slide with Kent Thiry, the CEO of DaVita. He thinks culture is very important. He makes it a big point of all his talks and all his corporate events. And I think that's what he's trying to get across, is that we want to reduce turnover because it affects our quality of care.

But another intriguing aspect of this industry, the independent facilities are often owned and operated by individual nephrologists, kidney doctors,

and they're being bought up, and they often leave right away. They're retiring or they just don't want to be part of a for-profit system. And, so to speak, that suggests that this is not a benevolent acquisition. They're really taking over facilities to increase DaVita's profits at the expense of potentially, you know, employee welfare as well.

MR. RASMUSEN: Hi, Eric Rasmusen, Indiana University.

Something neat about this is it looks like you can maybe point to the specific places where quality is going down and it matters. So you alluded to time between patients and septicemia, it sounds like could be caused only by that kind of unhooking and hooking up, or cleaning, whereas cardiac events is kind of vague. Maybe you can do more of that.

Also, what you find insignificantly different would be important in showing not -- where not to look for monitoring. And I wonder if you can see whether the cardiac events are mediated by septicemia, say, or not, so pin down exactly where the problem is.

MR. McDEVITT: I showed you a lot of results. There are certainly more results in the data. These are the most prominent ones, but certainly there's more scope for looking into some other measures of

quality.

I'll be frank about it, every quality measure we looked at, it's worse after acquisition, except for the couple that I alluded to here. We wanted to give a fair and balanced picture based on what we found, but everything we've looked at, patients are faring worse.

And there are clear channels, for instance, the cardiac events are clearly coming from the EPO doses that are too high. We can link those directly.

MR. LEWIS: So you motivated this as an issue of culture changing. So I'm wondering how much you can say about these effects being driven by just decreasing costs at the expense of taxpayers versus there's also some kind of transfer of culture, you know, via these acquisitions.

MR. McDEVITT: Yeah, I don't want to emphasize culture too much. I'm sorry if I gave that impression. What I meant to say is that culture is just an overall part of the firm strategy, and it's clear that strategy matters for these facilities.

And if DaVita, for instance, when they acquire a facility, they're transferring their strategy, culture is just one aspect of that. Probably the most direct example is that DaVita/Fresenius have extensive operating manuals they give to their facilities, over 100 pages, that tells you specifically what you should do in each case, what EPO dose you should have given a patient's blood levels, how long they should be on the machine, aspects like that.

So those are the types of strategies we're talking about, and we can see that borne out in the data, for instance, through the length of time they're on machines or the EPO doses. That's what we can observe.

MS. JIN: Can I ask a question? To what extent do you think consumers know those quality changes and still choose to stay versus they just don't observe those, or sort of it's so rare events that it sort of does not come back to them as quality and deterioration?

MR. McDEVITT: Hard to say how much information consumers have. Medicare makes available what they call a Dialysis Facility Compare Website. Very much like nursing homes, you can go to the website and see measures of how the facilities compare on infection rates, hospitalization rates, some of these measures. I don't know who's accessing this and if it matters, but what we find is that consumers are not responsive at all to quality.

Whether they know it or not, they just don't seem to switch facilities. Part of it is access. There needs to be an opening for them at a facility. Part of it is transportation costs, but we haven't looked at disentangling the information aspect, per se.

MS. MAJEWSKI: This is Sue Majewski from the Antitrust Division, Department of Justice.

I had a very related question, but typically we would be concerned about local markets and the acquisition's impact on a local market, and for that story to work, you have to have some sort of consumer substitution and some sort of signal why consumers would substitute a belief that they see some measure of quality, but I would love to see this paper sort of explore a local market angle with that.

MR. McDEVITT: As I mentioned, on the summaries, we're actually working on that as well. We're very intrigued by this. The preliminary results is there's just no response to consumers from local market concentration. And I'm going to rely, again, on this story of transportation costs. They just don't switch for whatever reason.

MR. GREENLEE: Patrick Greenlee also from the Antitrust Division.

Do you see any patterns in terms of capacity, whether it changes post-acquisition either in the number of machines per facility or upgrades to newer technology, if there are such things? Sort of in terms of capital equipment.

MR. McDEVITT: We don't have direct data on the actual machines they're using. We just have an overall count. The for-profit chains tend to have more machines per facility, and those go up a little bit after acquisition, but the issue for a facility and the standard is they just cram the facility with as many machines as they can, and then once they're at capacity, they build a new facility. It's really hard to keep adding machines.

Although independents may be a little subscale from a maximizing profit standpoint, but something we don't look into, which is another intriguing feature in this industry, Fresenius is vertically integrated into the machines. They're the main manufacturer of these. So another potential paper topic, and please don't steal it.

MS. DUTTA: I think we have time for another question.

MR. RAVAL: So given that it doesn't seem like consumers relate to local market competition, how

would you advise the regulator to change things to improve quality?

MR. McDEVITT: That's a big question and a fair one given where we are today. I think I would look also at the merging entities and have that as a part of antitrust regulation. It's not just at a local level, but it's what evidence we have of how these firms implement different strategies after acquisition, and in some ways discipline them on that, and have standards, for instance, how many patients you can have per station, employees per station, have some guidelines for EPO doses, more direct measures.

And that's probably outside the FTC/DOJ purview, but in this setting, that will be crucial, because competition doesn't seem to have much influence.

Thank you, everyone.

(Applause.)

(End of session.)

PAPER SESSION:
NONPARAMETRIC ESTIMATES OF DEMAND
IN THE CALIFORNIA HEALTH INSURANCE EXCHANGE

MS. DUTTA: All right, thanks, Ryan, and thanks everyone.

We're now going to move on to the second paper for this morning's session, which is titled Nonparametric Estimates of Demand in the California Health Insurance Exchange. I'm going to invite Pietro Tebaldi of the University of Chicago, who is one of the co-authors on this paper, to the podium to present it.

MR. TEBALDI: I also want to start by thanking the organizers. It's great to be here and to be on this program with this paper that is co-authored with Alex Torgovitsky and Habin Yang, who is now a student at Harvard Business School.

So as you can tell from the title, what we're looking at here is the context of the health insurance exchanges that were set up by the Affordable Care Act, Obamacare, if you want, in the jargon of the media, probably unnecessary in this room. This is a context where, as most of you know, we have consumers who are choosing a single insurance plan from a discrete -- a finite set of options.

We still have important policy questions that remain at least open in some sense, and examples of these questions are how would the demand respond if we're changing the premiums or the premium subsidies that the vast majority of these buyers are benefitting from in these exchanges, and what would be the corresponding change in consumer surplus?

Now, what we usually do -- or at least what I usually do in my other work -- is combining a bunch of functional form and distributional assumption that the random utility that is driving this discrete choice, where we think about the usual logit, nested logic, maybe multinomial probit if we can make it converge, or mixed logit, and then what we would ask in this paper is how are our results and maybe policy conclusions affected by these type of assumptions, and can we make important or informative conclusions avoiding these assumptions, okay?

So with this motivation, what we do in this paper is actually consider -- I mean, to give you like an overlook, right, what we usually have is this common practice, we have the parameterization of the random utility, we make assumption of how the unobservables and the coefficients are distributed, conditional on some observables (indiscernible), we

derive a likelihood, we estimate this model, and we end up with a point estimate that we can use to then extrapolate our demand curve at counterfactual prices.

What we do today is say, okay, we have the same model. We avoided this type of parametric and assumption on the indirect utility. We also consider a situation where we don't let us assume that we have these amazing instruments; we're moving the prices everywhere. We had a finite set of observed prices in the data. That means that what we're -- what we will end up is a partial identification framework where instead of a demand curve, I will end up with bounds on the demand curve, okay?

Importantly, we will show that these bounds are sharp, and a feature of this approach is that when we're going to ask a more ambitious question, which is like we move the prices in the counterfactual further away from the observed ones, the method with the wider bounds reflecting the higher uncertainty that we -- that we are facing as a researcher.

A second feature that I want to emphasize is that this method allow us to add assumption flexibly and in a transparent way, which is as you add the -- as you add stronger and stronger assumptions, this will tighten your bounds, and you can see as a

researcher exactly the role that is played by each of these assumptions.

I am going to guide you through the econometrics as intuitively and as parsimoniously as possible, and then I will show you how we apply our method to the context of the Health Insurance Exchange in California and how we end up with bounds on the demand changes and consumer surplus changes that are quite informative, okay?

Now, the model is a standard one. We have agents indexed by i. They make their changes from this set of j option. We have the prices that are collected in the usual vector Pi. We have a set of observables about the consumer, the market, or the goods that we collect in the vector Xi. I'm introducing these market indicators. That is an important piece of notation that I am not going to have a lot of time to emphasize today, but you'll want to think of this as the level at which your concern that the unobservables about these products or these markets are varying.

In the context of our application, this will be the rating region, which is the level at which the insurers are choosing how to enter in the exchanges, and they're setting the network of providers and their

45

premiums, okay? For today, we model these prices, these indicators, and these Xs as discrete. This is just to simplify this, what we want to show.

We put ourselves in the situation where the researcher is observing a collection of conditional choices. These are the standard market shares conditioned on what? Well, you will have how many people are buying j, given that they're facing the price b in market m and their Xs are x. We think of these as being possibly constructed from individual level data or observed already as market shares, but here, for today, I will treat this as given from the sky, go to the paper to think of the situation, as always, where we're actually estimating this maybe because we don't know exactly how many potential buyers we have in this market, okay?

The consumer problem, again, is standard, and this is where we introduce our main modeling assumption, which is that indirect utility is quasilinear in the numerator, in the price, or in the premium in our application, that the date equals zero is, as usual, the outside option, with the standard normalizations.

We don't impose any restriction on the joint dependence between these valuations for different

46

goods within individual, which is we avoid to impose restriction on the substitution patterns ex ante. Of course, we might be concerned that these VIs and the prices are dependent, and I will discuss how we want to think of instruments in this context.

This assumption that the indirect utility is quasilinear in the numerator is going to be the key for us to get traction here, which is, as you know, this means that now only the relative prices between two goods are going to matter. An implication is that if I increase all prices in the market by the same amount, I'm not going to change the relative shares between any pair of inside goods.

Now, this also means computationally that I can characterize the choice problem as a solution of a system of linear inequalities, and that is going to give us a massive computation and identification payoffs, as I'm going to show you in a couple of slides.

A byproduct of this model is also that we have a natural definition of consumer welfare because I can integrate over the VIs and I have a standard -- a standard utilitarian consumer surplus definition.

Now, in this context right there, the primitive object of interest, these are just a bunch of

47

definitions. I'm going to get to the meat very soon. In this context, right, the main unknown here is going to be the distribution of these valuations, which is unobserved, conditional on the prices, the market and the access. We are assuming that this distribution is harsh move and well behaved, which avoids us the ties and a bunch of irregularities using this problem.

Now, what it means now that we can work with this collection f calligraphic, if you want, that is going to be the collection of all possible conditional densities of these valuations, okay? And ideally, if we know this f, we can do anything, right, in this demand systems in terms of counterfactual.

Now, because of quasilinearity, if you give me a conditional density, I can derive easily the implied choices, right? Because as I said, like given quasilinearity, a consumer will choose good j if and only if his valuations are falling in the set. I can characterize this set as a system of linear inequalities. That means that the market share of good j at the given prices and observables is going to be the integral of this conditional density over this set that I can write down exactly.

Now, almost there. What we care in this context often, and in our paper for sure, is not the

48

entire distribution of these valuations, but instead, right, I usually have in mind a target parameter. I'm calling it Theta. This could be the change in a market share given a change in price, the change in consumer surplus given a change in price, and so on and so forth. All of these are functions of a conditional density, but they don't require to know the entire, density per se, unnecessarily, okay?

Now, imagine that I have this parameter in mind, as a researcher, okay? In our case, it will be the change in consumer surplus if I drop the premium subsidies in the ACA by $10 a month, okay? Have that in mind.

Now, I have this parameter. I'm going to make a bunch of assumptions on these densities, which is going to restrict the set of possible densities I want to consider. A standard assumption here will be a bunch of my Xs are excluded, are like a good instrument to think of exogenous variation in prices, okay, I can call this as a restriction on these conditional densities, and then, just by definition, what is identified here is, well, that the density that is consistent with both my assumptions and that is generating market shares that correspond to the observed ones, okay?

And now in terms of the parameter of interest, the sharp identified set for this parameter is going to be the image of this f* set under this function Theta, okay? This again is just a bunch of definitions, but it implies that if I was able to solve a very, very high-dimensional problem, and truly infinite to a dimensional problem, I could characterize the upper and lower bound on my parameter of interest as the solution of two problems, which is the mean and the max of what my parameter of interest, which could be, again, the change in consumer surplus in the California Exchange if I change the premium subsidies, over all the possible densities that are satisfying the assumption, and at the same time they generate the observed market shares. We know that this is true, but I personally and I don't think anyone here is able to solve these problems in practice because of the dimensionality.

So what our main idea of what we are trying to do here is take this and now say, well, I can rewrite this problem in a way that now is computationally tractable, and it gives me the identical solution to this problem on the top, okay?

Now, how does this work in practice? I'm going to show you this with one observed price and two

goods, and hopefully I can give you the main intuition, and then we go through the application that is perhaps more interesting.

So here we have the valuation of good one on the X axis, the valuation of good two on the Y axis. I put myself in the situation where we observe this price, pa, and we're interested in what? The counterfactual demand for good one, if I change the price from pa to p*, okay?

Under quasilinearity, I can partition the valuation space in three regions, the region of those who are buying good one in yellow, the region of those who choose the outside option in blue, and the region of those who buy good two in gray, okay?

Now, the observational equivalence means that I am only considering f conditional density valuations that are generating the observed market share in the data and they're integrating over these sets. I can do the same construction for the counterfactual price, okay, which means that my parameter of interest is actually the integral of f over the yellow region here, right, which is how many people will choose one if I am at p* and not at the observed price, pa.

This is the problem that ideally we want to solve, right? We want to maximize or minimize this

integral where the constraint is that the density is matching the observed market shares.

Now, how do I transform this problem in something that is tractable? Well, I'm going to consider this partition of the valuation space, that if you notice, what I'm doing here is intercepting these sets with these three sets that I had when I was considering the two prices either in the data or the relevant prices in the counterfactual.

So I'm considering this partition that has the following properties, right, that within each set, consumers are going to make the same choice at all of the prices that are relevant to this problem. Across two sets, consumers are going to make at least one different choice at either the prices we observe in the data or the counterfactual prices you care about in your research question.

But now I'm going to introduce the last piece of notation, which is I'm going to call "fee of L," the mass that the density of valuation is placing on each of these six sets. But now I can take my original problem and rewrite the objective as the sum of two integrals over the sets of the partition, and I can rewrite that the constraints also has some of integrals over the sets in the partition, and now I

can just plug in my notation, and I end up with this that is a finite linear program that I can solve easily. I know I have a unique solution, and I hope I convinced you that this is going to be identical to the solution of the infinite dimensional problem that I have up top.

And importantly, with engineering software, we can solve these problems with many, many thousands of parameters, or set of the partitions, if you want, very fast and efficiently, and we know that this is the unique solution to this problem, because of linearity, okay?

So this is what we do with two prices and no endogeneity issues and so on and so forth. In practice, in the paper, we go over all of the math that we need to extend this intuition to our general case, okay?

I'm going to skip through a little bit. I just want to say the instrument is something that we typically want to be concerned about, right, so here I was giving you the intuition in a world where the prices are exogenous. Now, this is not an attractive assumption. What we're going to do is kind of the standard thing here, we're assuming that a bunch of the covariates in the observables are going to be excluded or orthogonal from the valuations, and this is our IV assumptions, and notice that these can be encoded, again, as a set of linear constraints in that problem, and as long as your assumptions can be written as a bunch of linear inequalities or equalities, you are good to go, because you stay in the world of linear programming that we can trust very, very good software to give us answer very quickly. Okay?

I have six minutes, so I'm going to jump to the application to show you some numbers. What we are considering here is the California Exchange under the ACA, that they are familiar with this, so I'm going to move a bit quickly. We consider the subsidized population, which is those between 100 and 400 percent of the -- that fit our poverty level. We considered the choice between the four metal tiers in the market, and we have administrative data from the California Exchange, and here we are considering the first tier of the market, and different versions of the paper, we're adding the more recent years as well.

Now I'm going to jump to some figures. So in practice here, I'm showing you only the bronze and the silver because I can do it in a plane. These are the premiums that we observe in the data. The question is

what happens to demand and consumer surplus if, for example, I increase all of the bronze premiums by $10 a month, all of the silver premiums, or both of those premiums at the same time, which is equivalent to a reduction in the subsidies.

How do we think about identifying variation in this context? Well, in the ACA, after you tell me the region where you live, your household size, your age and your income, your premium is a deterministic function that is coming from the regulation, okay, which means now if I don't want to go across regions -- which I don't want to do because the unobservables are varying across rating regions -- to get variation in prices, I must extrapolate across households with similar characteristics.

So what we're doing in practice, we group our households in income bins that are in six income bins, as I'm showing here, and in age bins of five years or smaller. The assumption in our main estimates -- and then I will discuss how we can relax this somehow -- is that within each region and within the intersection of these income and age bins, the valuations have the same distribution, okay?

With this assumption, we can apply our method, and the first output is an elasticity matrix, if you

want, right? It's substitution patterns, and you see that instead of having a point in each entry of this matrix, I now have an interval, which is the sharp lower and upper bound on the substitution patterns that we estimate in this market. I just want to emphasize that if you look at this, these bounds are quite informative.

In particular, in the bottom right corner of this table, I can see that if we increase all of the premiums by $10 a month for all of the households in the market, the enrollment probability decreases between 3.3 and 8.4 percent. When you look at this table carefully, you also can notice that the substitution patterns do not expose IIA, which is like we see in the substitution between the bronze and the outside option is much higher than between higher tiers and the outside option, as you would expect.

If we look at different counterfactual prices, which is here, I'm showing you on the X axis, the change in premium for old plans in dollars per month, and on the Y axis, the probability of buying coverage, you see that our method, as I was mentioning earlier, is going to give you wider bounds as you extrapolate further away from the data, okay?

So what we know here is that the demand curve

is in between these two curves, but if I was to estimate and mix logit right, I would pick one inside or maybe outside of these intervals, okay? So that's kind of like the main output that we're getting here.

What we do then is consider consumer surplus and government spending. If I think about reducing all of the premiums or equivalently -- I'm sorry, increasing all of the premiums or equivalently reducing the premium subsidies by $10 a month, what we do here is minimize and maximize the area in this figure under the data constraints in our assumptions.

What we find is that, on aggregate -- I'm looking at the bottom row of this table -- you would save between 56 and 70 million dollars a year -- I'm sorry, that you would penalize consumers between 56 and 70 million dollars a year, but at the same time that you would save in government outlays between 440 million and 768 million dollars a year.

This, again, is the usual finding, that if you look at utilitarian consumer welfare in this context, you find that we are subsidizing people that don't value these goods too much. This is not a new finding, and we have a whole literature who's trying to explain why we see that these people, they don't seem to value health insurance as we would have in a

57

standard model, okay?

Now, I mentioned our assumption in terms of assuming that within these small age/income groups, the valuations don't vary. This is somewhat concerning, right, and what happens in other countries as well, you might want to relax your exclusional restriction and think of a situation in which you don't have a perfect instrument, but you might have different values of the instrument, your valuations are somewhat different.

Our approach allows to deal with this, and I just want to say this before I conclude, is, you know, you could think of a world where I don't want to say that for different ages the valuations are identical, but I am willing to take a bandwidth parameter Kappa and say, like, as you go from 31 to 32, your valuations don't change by more than 20 percent.

And, again, in this context, I can write this as a linear inequality, and I can run it through, and I can check the robustness of my estimates to the relaxation of the exclusion restriction, and I think I like this feature of what we're doing.

I am out of time. I am going to leave you with this figure, where I compare our estimated bounds to your standard parametric models. Maybe this is good

58

because they fall inside our bounds. One thing that we noticed and that we are trying to explore further is how we tend to kind of hit the lower end of the price sensitivity compared to what our model implies could be a worst case scenario in terms of demand responses to the premium changes.

And I'm totally out of time, so I'm going to leave you here. Thank you. Sorry.

(Applause.)

MS. DUTTA: Thank you, Pietro.

So, let me welcome Kate Ho of Princeton University to discuss the paper.

MS. HO: Thanks, and thanks to the organizers for putting together such a terrific conference. I've enjoyed it a lot.

So I enjoyed reading this paper. Let me sort of take a big step back and out of the details and think about context here, right? So if you think about the recent literature that estimates demand, particularly in medical care, consumer demand for hospitals or for health insurers, as Pietro said at the beginning, these models are often fully parametric, right? So we assume that consumer i chooses a plan or a hospital j based on its characteristics, we write down some indirect utility

59

equation that I have written out here, right? The utility from product j depends on its characteristics, that's x, and the price, that's p. We often have some unobserved quality x(c)(j) and we put on a logit error term, and we take choice data, either at the individual level or market shares at the product level, we estimate by maximum likelihood, and then policymakers go ahead and use those estimates when they're thinking about merger policy or regulating markets, right?

And I think there's a sort of healthy skepticism, or an understanding among practitioners that it's important to evaluate the robustness of those estimates to the assumptions that we're making along the way. So, you know, we use logit errors. What happens if we make some other assumption? Do things change if we put in brand fixed effects? if we put in interactions between consumer and product characteristics?

And many of you in this room know very well some examples of these kinds of papers. I stood here two years ago maybe and discussed this terrific paper written by people in the room that used natural disasters as an instrument essentially that unexpectedly remove hospitals from local markets as a

60

shock to help us evaluate these kinds of models, and that's a nice paper. I think it's R&R RAND? No? Yeah, okay.

So this is the context or one of the contexts that this paper can live in. This paper is trying to take a broader view. The authors say, well, okay, let's write down this consumer indirect utility equation in a slightly more general form, or arguably a considerably more general form. We don't want to make an assumption on parametric specification or distribution of these VIJs and see how far we can get with an essentially nonparametric model, right? We're only going to assume the valuations and premiums are additively separable.

So, of course, that relates to a large literature on semiparametric and nonparametric approaches to unordered discrete choice analysis dating back, way back to Manski in the seventies and Rosa Matzkin in the '90s, and these models inevitably are often partially identified. So I've written out just a few of those papers.

In thinking about sort of best practice and how to think about where to put this paper in that literature, I went back and read again a paper that I wrote actually fairly recently with Adam Rosen, trying

61

to provide a survey of this literature and suggestions for best practice. And it turns out that Pietro's paper checks many of the boxes, so I wanted to go through some of those boxes, right?

The paper essentially uses restrictions that are motivated from economic theory, right, with limited additional assumptions, and that's clearly a good thing. There's an idea that crops up in this literature that it might be sensible to try to place bounds directly on the counterfactual values of interest rather than on the underlying parameters, the betas in the utility equation.

Why might that make sense? Well, essentially because the values of interest it turns out are often much simpler and easier to bound than the underlying multidimensional parameters, and if you go straight to the counterfactual of interest, you might generate narrower bounds than if you go to the underlying utility equation and then inflate things up.

And clearly the authors are thinking hard about that kind of issue, and then their method provides sharp bounds. What does that mean? Well, it means that the bounds contain only parameter values that could have generated the data given the assumptions and no others. Well, that's clearly, you know, a

62

benefit in these kinds of approaches.

And finally, there's an idea that it's important when you're using these kinds of methods to explore the implications for the estimated bounds of relaxing the various assumptions you've made. That idea goes all the way back to Manski. Let's make very minimal assumptions and look at the bounds, and then let's layer on additional assumptions and see how much the bounds change. And the authors do some of that. I would suggest that they do more, so I'll come back to that a bit later on.

So, briefly, how does this method work? I'm going to take another stab at explaining what's going on here, because Pietro didn't have a ton of time, so let's see if I can make this make sense in two slides.

So here's the idea: Suppose consumers choose an insurance plan to maximize the indirect utility, right? We observe market shares of each product, given prices and Xs. The authors define what they call minimal relevant partitions. Remember that picture Pietro put up with the shaded regions, right, in different colors? These are minimal relevant partitions. They have sets of valuations that are observationally equivalent given the data.

So within each of these sets, all of the

63

valuations are going to generate the same observed shares for every vector of prices that we see or counterfactual vector of prices that's relevant, okay? And then we're going to move from the space of valuations v to the space of mass functions phi defined on those MRPs, right? Fine.

Then we're going to write down a familiar key condition, which is the predicted shares from the model equal observed shares in the data for every vector of observables, and, you know, that's fine; we do that all the time. And then we're going to notice that because we've moved from the space v, or the underlying parameters, to the space phi, this condition now generates simple linear constraints on the phis without throwing away any of the information in the data. That's kind of a clever idea, I think.

Notice a couple of things. More observed premium vectors provide more information, right? More observed premium vectors imply smaller sets of observationally equivalent valuations, hence more MRPs and more linear constraints. And so intuitively, the more prices we observe, the narrower the bounds are going to be, and that makes sense.

And then we can add further conditions, instruments and a vertical assumption that I won't

64

talk about in detail, and then finally, we're going to define our objective interest. They call it a target parameter theta, preferably as a linear function of these phi's, and the phi's then have to be sufficiently rich in order to fully determine the target parameter of interest, the change in consumer surplus, or a change in market shares with a change in policy.

And then we can place bounds on this theta of phi, just at the lowest and highest values, such that all of the linear conditions on phi is satisfied, and notice that that's a linear programming problem. There may be thousands of constraints, but still it's relatively simple and it's going to generate sharp bounds. So that's the idea in two slides. I think it's a very nice method.

So some of the ideas here, of course, go back to themes that are dispersed through the literature, but then a lot of them are new and pretty creative. The authors say in the paper that many previous partially identified models deal with the unobserved components of indirect utility, sort of the epsilons or the Cs or these components of the v(i)(j), but they deal with them as a nuisance parameter, and a lot of my own previous work does this. They don't try to directly estimate them. They're just trying to deal with them.

This method doesn't do that, and the authors, I think quite rightly, point out that that's a benefit, because policy counterfactuals also depend on the distribution of these unobservables. So that's nice.

The method allows prices to be endogenous if you can come up with instruments that you believe, so that's obviously a benefit. And, by the way, the authors note that relaxing the usual point identification assumptions may matter, does seem to matter, for policy relevant objects like the effect of this premium subsidy change on consumer surplus. So that's potentially important.

Okay, a couple of specific questions and comments. So my sort of overall comment is that I like the method. This paper was a tough read. You know, if I were you, I would focus on making this useful for empirical people, for practitioners. When I read the paper for -- you know, I read it several times, and even on the last reading, to me there's a disconnect between Section 3, the method, the econometrics, and Section 4, the empirical application, right? Some of the details of exactly what you're really doing are in the appendix. Some of

them are not anywhere, I don't think.

If I were you, I would take this estimation section, which I had to read through to Appendix G to find, put it in the paper, right, use up a page of text and just lay out a menu for practitioners exactly what are the equations you're using for estimation. Put in more explanation.

Some of these results are super-important, I think, about the impact of these policy changes on government spending, for example, and I needed more explanation of exactly how you got there. And then you can discuss how to generalize, relaxing the IV assumption, which you do very nicely.

Can you allow for horizontal differentiation? What if you put in brand effects, how would that change the method? And the kinds of issues that practitioners are going to care about.

And then sort of thinking back to my slide on best practice, I think it's super-important in this literature to explore the implications of relaxing the assumptions for the breadth of the estimated bounds. And certainly there's some of that in the paper, but I would have liked to have seen more.

So what happens if the vertical structure constraints are removed or relaxed? It's pretty clear

that the IV assumptions are crucial for estimation, and you do a nice job of relaxing them, but I didn't see -- I don't think I saw what would happen if you removed them entirely. Perhaps that's in the table, but, you know, in the Manski framework, starting very broad and moving inwards rather than starting in and moving out I think would have been useful.

One more slide and then I'll be out of time. So overall I really think this is a creative and intuitive idea. This idea of redefining the objective interest in terms of objects, these phis where there are linear constraints still generating sharp bounds, everything gets much simpler once we're in a linear world.

It seems to me that there are two tricky steps here, right? The first is characterizing these sets, these MRPs. Even in the case with only two products, you know, the picture looked a little bit complicated when there are many products and we've got instruments, and I'm sure it's an extremely involved process.

And then secondly -- and these two things seem to me to be intertwined -- that the challenge of defining a target parameter or an object of interest that's preferably a linear function of these phis,

right? And it would be great if you could say something more about broader applications with these challenges in mind.

So what other target parameters could be assessed using this method? Did you have to very carefully pick this target parameter in order for it to fit into the methodology? It seems to me there's a tradeoff between, you know, defining an interesting counterfactual versus needing a large number of MRPs in order to determine it and that part of the process being implausible. I may be wrong about that, but it would be great to see more discussion of that issue.

And then going back to what I said on the previous slide, so you're characterizing sharp bounds, and in the results that are also informative, exactly which assumption is it that makes them informative? And as you change the assumptions, how much do the bounds change?

But overall I think it's a terrific paper. I think it's going to have some real impact. Really, my only comment is make it easier for us to understand so that we can actually use it.

Thanks a lot.

(Applause.)

MS. DUTTA: All right. Thank you, Kate.

So we can have ten minutes again for Q&A.

AUDIENCE MEMBER: So I think the assumption that the functional Theta is linear, you need to have that to get a connected, sharp identified set. Is that where you're actually using it, or is it bringing you more?

And related to that, what kind of -- with this linearity assumption, what kind of target parameters are you actually excluding?

MR. TEBALDI: Sorry. So the first one.

AUDIENCE MEMBER: So the first one --

MR. TEBALDI: No, no, no, I remember. So the first part is do you need a linearity for the sharpness, the answer is no.

AUDIENCE MEMBER: For the connectedness of the sharp identified set.

MR. TEBALDI: So you can look at it -- so if you look in the paper, what we show is first how you transform the problem in the finite problem. In that problem, you know that the problem is regular, and it gives you a connected set.

Now, however, what we cannot -- so the key here is that to transform the problem into the finite problem, you need functions that only vary with the mass over the MRP, okay? And that's the key here

because -- I mean, obviously, like, if I transform the problem in a problem that only depends on the mass, the objective function needs only to depend on this mass.

That restricts us -- and if you think about it, if you ask a question that depends on anything more than that, which is the heterogeneity within these sets, there is nothing in the data that can possibly tell you that, right, because the definition of the partition is such that all the information that you have in the data is contained in this problem, which is -- in some sense you're limiting yourself to target parameters where you possibly have information in the data.

Now, what this allowed us to do is the demand stuff that I showed, I saw like the counterfactual choice shares and things alike, the consumer surplus changes, and similar other problems. What it doesn't allow us to do is anything that has to do with integrating within the sets of the partition, and we can come up with a lot of these questions of interest.

AUDIENCE MEMBER: I guess I just didn't understand the requirement that the target parameter is linear and the -- the one phi I guess.

MR. TEBALDI: Linearity means that I can write

it as an integral over sets of the partition, and that means that if I'm integrating over a specific function, I need to assume that this function is known ex ante. And like in the case of the market share, it's an indicator function; in the case of the consumer surplus, we know how to write it down. That's kind of the answer.

MR. LEWIS: Is it possible to pull up your slides?

MR. TEBALDI: I don't know. I have very little control.

MR. LEWIS: So I was wondering if we can get the slides up, and if you can show that graph of the partitions, and just walk through the intuition about -- so the intuition of you had the division of, you know, don't buy anything, buy good one or buy good two.

MR. TEBALDI: Yeah.

MR. LEWIS: And then you had the further partitions, and I was wondering if you could just walk through that and explain how that relates to the consumer surplus question.

MR. TEBALDI: Oh, how do we deal with the consumer surplus question?

MR. LEWIS: So you're using this partition to

specifically get at the issue of what is this target parameter of interest.

MR. TEBALDI: Yeah, yeah, yeah. So actually it's related to --

MR. LEWIS: And so how did you draw those partitions such that that answers that question?

MR. TEBALDI: So the partition is not -- okay, so let me go here, right? So the partition really doesn't depend on the question. The partition is going to depend on the prices you observe in the data, in this case only one, Pa, in blue, and the prices that you care about in the counterfactual, which is p*, okay?

And this is -- the partition only depends on the prices you consider. If you give me a set of prices, I end up with this, which is as I cross between two sets, at least one of these price's agents are going to make a different choice, okay? And if I stay within a set, at all of the prices that are relevant to this problem, consumers are making the same choices, okay?

MR. LEWIS: Right. So the idea is -- so the V2, for example, is saying that both in the world with pa as well as in the world of p*, people who are in that V2 area are making the same choice regardless?

73

MR. TEBALDI: Yeah, and they are going to choose good one at pa and they're going to switch the outside option at p*. These guys in V1, they're going to stay in the outside option at both prices, and we can go over all of these.

MR. LEWIS: I see.

MR. TEBALDI: Okay? Now, for the consumer surplus, you realize that when we go from pa to p*, we should take the integral over all of these Vs within V4 maybe, okay? And this is answering maybe your question, but that's where you go back to the drawing board and you realize that to characterize the upper and lower bound of consumer surplus, within each set of the partition, I can place all of the mass at the extreme points at either the southwest or the northeast point for the consumer surplus. For the demand changes, I don't need to do any of that. It's kind of hard without the slide.

AUDIENCE MEMBER: Pietro, thanks. I was going to ask how you deal with derivatives, because all these are integrals, but I guess when you have enough price changes, these little phis are going to be over small enough grid points that you're going to look at -- like elasticities, right? These are derivatives.

74

MR. TEBALDI: Yeah.

AUDIENCE MEMBER: What you're identifying are integrals or -- but I guess you're looking at small changes, and if you have enough prices, with small enough cells, that's how you do -- you know, elasticities?

MR. TEBALDI: Oh, so actually what we do here is change in shares, and you're right that we often -- we want to transform these in the same elasticity or elasticities. So for the -- we didn't think about this idea of observing a small change in price, and we could do that if you want to then interpret it as an approximation of an elasticity.

What we did think about and we are trying to work it through is there is a -- this part of linear programming that is called fractional linear programming, where even if you don't have a linear objective function, that you can deal with ratios and keep the sharpness. This is something that goes a little bit beyond my contribution to the paper, but you can deal with same elasticity and elasticity in this formal way, but it's a good idea that if I had small variation in prices, I could do this as an approximation, just directly using our method.

MR. DUTTA: So I think we're out of time. So

75

thanks, everyone. We're going to take a short break and be back here at 11:00 for the keynote address.

(Applause.)

(End of session.)

76

KEYNOTE ADDRESS, "OWNERSHIP CONCENTRATION AND STRATEGIC SUPPLY REDUCTION"

MS. DUTTA: All right. Welcome back, everyone.

So I have the great pleasure of introducing this morning's keynote speaker, Dr. Katja Seim, who is an Associate Professor of Business Economics and Public Policy at the Wharton School. She is also a member of our scientific committee this year.

Dr. Seim's research has been published in a number of leading journals, including American Economic Review, the RAND Journal of Economics and American Economic Journal-Microeconomics. Her research primarily focuses on two broad areas in empirical industrial organization, which are product introduction and entry decisions by firms, and price discrimination and nonlinear pricing, especially in the context of the communications and information industries.

In addition to her research and academic work, Dr. Seim has also served as the Chief Economist at the Federal Communications Commission. Dr. Seim's keynote address today is titled Ownership Concentration and Strategic Supply Reduction. Please join me in welcoming Dr. Katja Seim.

(Applause.)

MS. SEIM: All right. Thank you very much for having me, including me on the program and asking me to contribute. Oftentimes I think when you're asked to give a keynote address, all of a sudden you feel very old, and I wasn't quite prepared to feel so old. So I decided to use my time today to actually talk about a project of mine that I have been working on again pretty actively rather than giving you a more general talk about the state of the literature.

This is a paper we've worked on for a while and are about to hopefully wrap up again pretty soon. So suggestions and comments would be very appreciated. It's joint with a number of people who always were colleagues of mine at Penn but have mostly moved on since, Ulrich Doraszelski, Mike Sinkinson and Peichun Wang. The paper broadly looks at the role of market power by TV broadcast stations in the recently completed spectrum auction that the FCC ran that is called the incentive auction.

And so before telling you a little bit about the details of what we do, I also wanted to say that I was very fortunate to spend some time at the FCC, like Antara said, during the time that we spent on this project, and so I should probably actually have the disclaimer that other government people have on their

slides, but most importantly, I think for our purposes, it's really been helpful in helping us understand the auction and being able to draw on experts who understand it even today, I think, much, much better than we do. And so I was really grateful for all of the support we got there.

So in motivation, let me just talk a little bit about the types of issues the FCC faced and what they attempted to do. As I think we can all agree, the biggest resource, the FCC allocates its spectrum, which is scarce and maybe increasingly valuable, and one problem that then arises is that we have lots of allocations of spectrums for two types of services that historically were very intensively used but are maybe less intensively used today.

The biggest such service is broadcast television, which is over-the-air TV. Today probably less than 10 percent of households use over-the-air TV to access programming. Nevertheless, as this slide shows you, significant amounts of spectrum, even today, are dedicated to broadcast TV transmission.

And so what you see here in this slide is basically two portions of the band plan that is set aside for broadcast television. The lower portion, which is allocated to VHF, low intensity services, we

will not talk much about. We are going to focus on the upper portion of the band plan, which is allocated to UHF TV channels. That has shrunk over time because of the digital transmission introduction by these stations, but even as of today, we have set aside spectrum for 37 different UHF channels in several markets, and that spectrum occupies space that we might think could be more efficiently used by other service providers, in particular cellular and wireless providers.

And so what I have shown you, then, at the bottom is what the band plan hopefully will look like by 2020 when the amount of spectrum to UHF channels has shrunk from 37 to 23 channels, with the remainder of the spectrum moving to wireless providers.

And so, you know, what we look at in this paper is how might you facilitate this kind of transition of spectrum allocation and to what extent do individual firms that own multiple broadcast TV licenses interfere with how that process works.

The challenge that we face in facilitating that kind of transition is twofold. The one on the TV side spectrum is very fragmented. In ownership, there's about 2000 full-power TV stations today that have market areas that did not line up at all with the

wireless market areas, and so here in this slide you can see the contour of ABC New York and the market area that it serves, that has very little to do with the type of market that a wireless provider typically thinks about.

At the same time, then, what we want to think about is, you know, in allocating a spectrum, how we incentivize TV broadcasters to give up their spectrum voluntarily and re-allocate it to wireless users who will hopefully use it more efficiently. The FCC decided between 2016 and '17 to run what they called an incentive auction. It's really unique in a number of ways. Most importantly, it's the first auction that they have run where they both bought up spectrum and then turned around and resold that spectrum at the same time.

The paper here is going to focus on that initial buy-up process, which is called a reverse auction. The second portion of it of then turning around and selling the spectrum in the mobile market is the forward auction. These two interact in that the biggest challenge the FCC faced in the auction is that it needed to reduce the amount of spectrum allocated to TV broadcasters before turning it over to mobile broadband service and do what they call repacking stations into a smaller portion of the spectrum band plan.

So, you know, the way this worked -- and, you know, if you take away one thing from the presentation, I think it is that this is an extremely complex process that I think the FCC deserves a lot of credit for having pulled off in a very smoothly running auction. What needs to happen is that a station who currently broadcasts on, say, channel 45, and does not want to give up its spectrum in the auction, needs to be artificially moved down to the portion of the spectrum that continues being TV broadcast spectrum, so say channel 27.

The FCC does have the right to move stations, even though they don't have the right to force stations to give up their spectrum, but they can only move a station as long as that move limits the amount of additional interference that the station faces to less than 0.5 percent of its current ownership.

And so, you know, that basically means that I can move you to a channel as long as the amount of population that you are serving after that move is extremely similar to what you were serving before. So that is going to introduce a bunch of constraints on who can be located next to each other in the spectrum

plan.

And so the biggest challenge, then, with the incentive auction is that what their goal was not only to identify the lowest cost set of stations that they wanted to purchase to acquire a certain amount of spectrum, but it was a constraint problem in that whoever chose not to sell out needed to be repacked into the remaining channels.

And so if you then think about, you know, this repacking problem, there's many combinations of stations that could potentially remain on the air after the auction, and that in terms of a feasibility checker is the biggest computational issue with the auction. It means that every stage of the auction computationally checked that whoever is left can still continue broadcasting.

The way the auction worked is you should think of it as a descending clock auction. It was developed by Milgrom and Segal and was operationalized through a nationwide base clock price that we've called capital P here, but that clock price was translated into an individualized price for every station as a function of what they called the station's broadcast volume fee.

The station's broadcast volume fee simply

reflected two things that resulted in stations being differentiated. One, it reflected the station's population reach as a measure of how attractive the station was to viewers. And then two, it reflected how difficult the station was to be repacked should it choose to stay on the air and that they proxied for by the number of other stations that the station could not interfere with were it to stay on the air.

So these interference constraints basically tell you, you know, if you stay on, how difficult is it for us to fit you into that remaining amount of spectrum, and as a result, if you're really difficult to fit, we want to incentivize you to actually sell out. And the broadcast volume reflects that.

So then think about the strategy. The nicest thing about the Milgrom and Segal descending clock auction format is that if you own a single license in that setup, it's weakly your dominant strategy to bid your value. In our context here, because of this relationship between the nationwide base clock price and your broadcast volume, that basically means you stay in the auction until the nationwide clock price drops below your valuation adjusted by your broadcast volume.

And everybody follows that strategy in

equilibrium, and we can, therefore, then identify how expensive it would be to buy up a certain amount of spectrum that would then be turned around to the forward auction, and we can figure out what would wireless companies be willing to sell to purchase that spectrum. And if their willingness to pay does not exceed what the broadcasters need to get to sell off that spectrum, we would lower the amount of spectrum and continue; otherwise the auction is going to close.

Now, this works nicely if there's single owners, every station is owned by a single company. What we're interested in in the project is what the role of multi-owners might be who own basically chains of broadcasters who own several stations and might have the ability to strategically interfere with how efficiently the auction works.

There's multi-license ownership for two reasons. One is purely historical accident. The FCC is typically concerned about market power in the advertising market in a local market, so there's constraints on which types of stations you can own jointly. You can't own ABC and CBS, for example, but there's typically much less concern about you owning multiple independent stations that are not very valuable, but are valuable from a spectrum perspective

85

if you were having concentration.

How we got interested in the project is that after the announcement of the auction, we also saw a lot of buyouts of largely failing stations by private equity firms, which amassed significant spectrum holdings, and I've just here shown you a bunch of listings for one of the three companies that were concerned, NRGTV.

So you can see that they bought out a whole bunch of stations, most of them are on the coast, not particularly successful stations from a broadcast perspective, but potentially quite valuable from a spectrum perspective. This drew a lot of attention from a speculative perspective and resulted in, you know, for example, bidding wars for what were nearly bankrupt stations.

The trade press focused on flipping of these stations, which isn't necessarily an efficiency problem. We're also going to think about the possible role that strategically you might have from owning multiple stations. So this gives you sort of an overview of these three private equity firms which attracted a lot of the attention, this NRJ, OTA and Locust Point. They bought up, up until the onset of the auctions, about 44 licenses, but you should keep

86

in mind that this is only a small fraction of the multi-license ownership that we see, the remainder largely being that way before the auction even was announced.

And so what we then want to think about in the paper is what the incentives might be for these firms to strategically withhold licenses from the auction. We are going to think about, you know, your ability to affect the base clock price and the price at which you might sell other licenses that you own as potentially affected by your decision to bid or not bid in all of the licenses that you have.

And so this is going to be similar to the types of strategic supply reduction effects studied in the electricity markets. So Ali, for example has done some work there, but I want you to keep in mind, in terms of how we're different, it's basically that A units are discrete, and we don't see firms bidding supply schedules, but maybe more importantly, TV stations are not a homogenous product because of the way they interfere with each other, which isn't necessarily the case in electricity markets, and so it's a nice setting to think about what the role of product differentiation might be on your ability to move prices in these types of settings.

87

And so we do three things in the paper, and I don't think I will really talk about very much of any of these, but we first develop a simple model of the reverse auction that gives you a sense of when you might want to successfully withdraw a license from participation, to think about, you know, how you can exploit the fact that you own multiple stations, but then going to try and quantify the possible contribution of such behavior to the auction or auction's outcome by first estimating reservation prices for all of the participation stations, and then quantifying the impact of strategic bidding using a large-scale simulation of what would have happened in the auction had such strategic behavior been important.

So we are going to do this both for the initial clearing target that the FCC announced of freeing up 126 megahertz of spectrum and the one that was then finally realized of freeing up 84 megahertz. So to give you a little bit of intuition on the model side, I wanted to make sure you understand how repacking works and how that affects our findings. Let me move on to that first.

And so to think about how repacking works and how it interacts with strategic behavior by firms,

88

we're going to think about clearing 126 megahertz of spectrum, going to leave us about 16 channels that stations who do not want to sell out can move to. A stations's decision if they bid naively based on just single individual station behavior is they're going to stay until the clock price falls below their value as operating a TV station.

If the station withdraws and chooses not to sell, it then needs to be repacked into that lower portion of the spectrum. The system is going to verify that all of the remaining stations could also repack should they choose to withdraw later in the auction. And if a station cannot be repacked, it would then be labeled as a winner of the auction.

This is going to continue until all of the licenses are repacked or are winning, and the auction is then going to conclude. We give you an example of Philadelphia. We like Philadelphia for a number of reasons, one of them being that strategic supply reduction is really only important in large markets where wireless demand is high. Philadelphia is one of them.

And so what I'm showing you here is interference constraints for a station in central Philadelphia, NBC Philly. In that, we show you

89

adjacent channel constraints, stations that cannot be located right next to Philadelphia in the channel lineup, and then in yellow we have stations that cannot be located on the same channel.

So this is where we start. For example, today, we have here the current set of channels before the auction. We've shaded in light blue the ones that are currently occupied by stations, each of which have six megahertz of spectrum, and what we want to do is take these stations -- some of them are going to go off the air, some of them want to continue -- and squish them into that smaller portion of the spectrum.

So, for example, the clock price is going to tick down, tick down, starting at 900, and initially not everybody stays in the auction. The price is high enough that they would prefer to take it as opposed to leaving. Then, say, we hit a clock price of 600. At that point, the first station withdraws. In Philly, the most valuable is CBS Philadelphia. They choose to exit the auction and continue as a broadcast station.

At that point, when only CBS is there, everybody else is currently still active, could still fit in that smaller portion of the band plan, the clock ticks down to 550. Now NBC goes out of the auction. Everybody else can still be repacked. Fox

90

drops out at 500, and then we're going to hit 450. MyNet is out. Everybody else can still be repacked.

Now we hit 400, and when we hit 400 and Univision decides that they would prefer to continue operating, holding onto their license rather than taking that price, is actually going to go out, but at this point there are three stations that -- with the existing five stations that continue operating could no longer fit in that band plan.

So at this point, the FCC is going to deem these stations as conditionally winning, and they would get, in terms of price, the price at which Univision, the last station, would go out, so it's a clock price of 400.

And then the process continues, and we're going to keep going until all of the stations either have chosen to exit the auction, because the clock price has fallen too much, or can no longer be repacked, in which case they are conditionally winning.

So, you know, then think about how strategic supply interaction works here. If you are a single license owner, you just follow your dominant strategy, and it's a second price auction where the price that you get is set by the firm that leaves just before you.

91

For a multi-license owner, we are going to think about what would happen if you considered the following strategy. You have, say, two licenses. You can consider to take one out and simply bid the other one at the naive strategy off the value, in which case withdrawing the first one can help you for total payoffs for two reasons.

One, it might mean that your other station no longer can be repacked and becomes a winner at a higher price; or two, it could be that by you not participating, there's a different station that sets your station's price that might also be better.

We developed the theory in the paper a little bit. None of us are auction theorists, so that's one small problem, but we can show that this strategy of either withdrawing a station and continuing to bid at value for the other one that you have is a weakly dominant strategy for these strategic bidders.

So let me just give you an example and then I'll show you our results. So think about you being a TV broadcaster that owns two licenses in a market, a and b. These licenses are in a market where the FCC plans to buy K licenses, and we've ordered the licenses by the drop-out point, which is the value adjusted by the broadcast volume.

92

Now think about the case where both of these two stations that you own have, in terms of their score value of a broadcast license, a score such that they are within the set of K stations that should be bought. And under naive bidding, where everybody bids their value, the K+ first station would set the price for all of them.

In that case, the firm's profits is going to be what it gets back in the auction, which is the base clock price of the K+ first station, scared up by broadcast volume, minus the value that they give up the profit from being a broadcast station.

In contrast, if they decided to withdraw one of the licenses, that might give them higher payoffs because we've now raised the closing clock price to be the score of the K+ second station, and our firm is now going to, in terms of outcomes, make profit on the late license b, which it sells, and hold onto license a, which has outside payoff value.

And so that's going to be profitable if the payout increases on license b, which it continues to sell, exceed the cost of no longer selling station A. And so if you think about then what types of stations qualify there, the opportunity costs of not getting rid of one of your stations is going to be low if that station is very profitable in the broadcast market -- that's what we call VA -- or if it has relatively low broadcast volume and can be repacked easily in the auction. And so that's sort of the types of tradeoffs that we evaluate in this simulation.

I should say this is a very complicated problem and with a nationwide system totally infeasible, so we do not attempt to make any such calculations. Instead, what we're going to consider is a situation that multi-license owners are strategic only in the market in which their license is, so in one DMA, and then we're going to also simplify this repacking system to be only applicable to the DMA itself and an area around it.

And so I sort of have some numbers that hopefully suggest that this is computationally complicated and tell you why we've been working on this for a long time.

Now, we're going to estimate reservation values -- I won't have much to say to that -- to be then able to quantify the value of strategic bidding. We basically use a cash flow model that is used by the industry as well to think about what the station's value going into the auction would be and what they would be giving up were they to surrender their

license.

One thing I should say is that this is unusual in the auction literature, that typically we rely on the model to give us optimality conditions, and we use that to back out the value that the station would have to have to be consistent with the actions that they take. We choose to do the reverse and start with an estimate of reservation values, and probably not a very good one, but we do that for the following reason.

In our case, the only data that we have about the auction and its outcome is the set of stations that sold and the price at which they sold. We do not know who participated and we do not know what the bidding behavior would be of a station that was frozen and how much further they would have been willing to stay in the auction had they not been frozen out.

And so the data that we have we feel is not particularly informative, even though I want you to keep that in mind when you think about what we do then in terms of estimating reservation prices directly.

So let me show you an example from Philadelphia and then the final overall results to illustrate what we do. This is our estimates of reservation values in Philadelphia. And maybe not surprisingly, they line

up pretty closely with advertising revenue that these stations can make. That's their main source of revenue and their main source of profit.

I want you to see that for the most part we see really skewed distributions. There are some stations, like the big ABCs, CBSs of the world that are very valuable. There's typically a large tail of stations that have very little value in the broadcast market.

Okay. Now looking at naive bids, we've overlaid your naive bids here, how long would you stay in the auction if you just bid your valuation. You can see that that also lines up nicely with the values but departs sometimes because your value in the auction also reflects your broadcast volume and not just your reservation value. And so that is, in particular, important for some of the low-value stations that interfere with a lot of others to actually be quite valuable in the auction itself.

Then we want to look at what that would look like under strategic behavior. We're going to simulate strategic outcomes using some of the FCC's own software to check whether stations can be repacked, and then use that to compare payouts under strategic bidding to payouts under naive bidding.

So, you know, we're going to repack the DMA

neighborhood like that, and in Philadelphia, then, this is what this is going to look like. Starting with the same chart as before, we have reservation values in blue, the naive bids as crosses, and then in dark blue the payouts that the stations received that were able to sell.

Now looking at how that changes under strategic, what I want you to focus on is two sets of multi-license owners in Philadelphia, and we're going to think about their payoffs should these two stations withhold one station each from the auction, at which point now their bids would be the same as before, except for they would ramp up the bids on these stations that they withdrew to just not be in the auction at all.

And so then what I've overlaid here is their payoffs under strategic bidding, and there's two things I want you to see. Strategic bidding here benefits the individual stations that withdraw their licenses from the auction, so that's the first thing. That just means that for them, it was individually profit-maximizing to do that. But they also impose a large externality on the other stations in the market that are single owners in that you can see that payouts increased across the board for the stations that sold. And so one reason why we find big payout increases is simply because there's this externality that strategic behavior benefits not just the firm itself, but everybody.

So I would show you the main results, and then I wanted to talk a little bit about how this compares to the actual auction outcome. There's a lot of numbers here, not all of which are important for what I wanted to show you. The main numbers that we take away from this is if we compare this naive bidding -- the everybody bids value -- to strategic bidding, we see that under the initial large clearing target of 126 megahertz, payouts would increase by 22 percent.

Under the smaller clearing target that was ultimately realized of 84 megahertz, we still see that strategic behavior increases payouts by 7 percent to the firms, and that is true across the board, both for single and for multi-license owners.

There's a number of caveats that we look at in the paper. There's two that are important. The first one is in our simulation so far, we've assumed that everybody participates in the auction. In practice, there were significant concerns about whether especially religious and nonprofit stations would actually choose to participate, and so we've redone

this under reduced participation and find, maybe not surprisingly, that the payout increases from strategic bidding go up very significantly, both for the 126 clearing target and the 84-megahertz clearing target.

The second one that I just showed you a map to illustrate how this could be a problem is we've considered strategic bidding in a DMA only, but oftentimes there's interference across DMAs that are nearby, and so you can see here, this is two stations that are owned by the same company, by NRJ. They are in adjacent markets, and we find that if they were able to bid those in both strategically, we would actually see pretty significant effects.

And so, you know, here we find, again, about 90 percent payout increases from this particular larger strategic bidding area. And so I wanted to tell you this, in just putting the auction results themselves into perspective, and that's where I'll stop. The auction, as I already told you, actually ended up geolocating 84 megahertz to wireless companies. It started with a clearing target of 126.

At that point, the demands in the reverse auction far exceeded the willingness to pay by wireless providers, and so there were three subsequent stages where those two were brought more into balance.

The auction concluded with 10 billion in reverse auction costs, more than matched by about 20 billion of revenue that the wireless providers were willing to pay for that spectrum in the forward auction. And so at that point the auction concluded, and we're hoping that by mid-2020, all of that repurposing will have been realized.

So, you know, in reconciliation, then, just to maybe remind you of our numbers, right, in our numbers we estimated that the true value to firms at the initial clearing target of 126 was not $86 billion, but just much less than that, about five times less. And then similarly, in our simulations, we also don't find anything near the 10 billion that was finally realized at the 84-megahertz clearing target.

And so in terms of why we were not able to match that at all, I wanted to remind you of these very conservative estimates that we provide by thinking about full participation and limiting strategic bidding to the MSA/DMA itself that you are in, and our results suggest that increasing those or relaxing those constraints would have given you significantly higher payouts than what we find here.

What we're currently working on a little bit is trying to think about, you know, what happened after

the auction. Some of our speculators did not, in fact, sell all of the stations that they had, which is one piece of evidence that we had that for them maybe this was not just flipping, and so you have sort of some numbers here.

And, you know, in conclusion, I would just say, with the data at hand, it's hard for us to prove conclusively that such strategic behavior was in effect, but one thing we wanted to point out with our work, since repurposing the spectrum going forward is similarly a problem, is that market power in these auctions, to the extent that firms realize that they have it, can actually have pretty significant effects.

And, you know, what we're currently then thinking about is, well, if we think about firms being differentiated, how would that change once we relax how much they can interfere with each other and, therefore, become more or less substitutable in the auction.

So that's all I had. Thank you very much for your attention, and I think we have maybe like two minutes for questions, since I went over my little allotted time.

AUDIENCE MEMBER: I have two questions. One, so maybe the other people here wouldn't be willing or don't want to point this out, but, you know, this is -- there was a campaign of buying multiple licenses by a number of speculators, or however you want to characterize them, in advance of this auction. You know, to my mind, acquiring assets in order to withdraw them or potentially withdraw them from a public auction would seem to be a kind of textbook violation of Section 7 of the Clayton Act. I wonder if you wanted to speak to that.

And then secondly, you know, I've heard a lot -- I have never worked at the FCC, as you have, about that forthcoming spectrum re-allocation. Is there any -- do you know of any sort of initiative there to treat the issues that you've brought up with the last one?

MS. SEIM: Let me just maybe answer the second one. I don't think right now there's any efforts in place to think about renewed spectrum repurposing. In part, I think that is -- you know, I think the big takeaway of the auction is running a complex auction like this is a difficult undertaking that has taken a significant amount of time. And so, you know, if you think about sort of what I showed you there on the slide before about reconciliation and what the expected bidders might have been compared to what they

ended up being, demand in these markets changes pretty dramatically, right, and one thing that I think has happened was that over the time period that we took to develop the auction, wireless demand has changed a lot, and as a result, you know, that spectrum wasn't as valuable anymore as it might have been when the auction was initially conceived.

And so I think as a result maybe of that, even though I think the auction ran very efficiently and obviously repurposed I think a lot of spectrum, there's less appetite maybe at this current stage to think about doing something like this again. And so to that extent, I'm sort of unable to answer your question directly.

Now, as for the antitrust question, there's a lot of people who have asked us that and who think we should make that the hangup of the paper. We're, I think, a little less willing to go there, mostly because we feel that we don't have any clear evidence that this was the strategy of these companies going in. Flipping, per se, isn't -- I think in any mean, way, shape an antitrust violation, and so to us we feel like we would need to have more evidence that we could clearly point to that that was something they had actually wanted to pursue. And I think you, you

know, can all appreciate that kind of evidence, while we can provide suggestive evidence from our results, it's a bit hard to come by.

All right, thank you very much.

(Applause.)

MR. ROSENBAUM: Thank you very much, Katja.

(End of session.)

PANEL: ESTIMATING MARKUPS

MR. ROSENBAUM: It's my pleasure to turn the program over to my colleague, Devesh Raval, who will be moderating a panel on estimating markups to conclude the conference.

MR. RAVAL: All right. Thank you all for coming here. So this is a panel on estimating markups, and we're privileged to have a great list of panelists here to talk about this.

So the first is Ariel Pakes, who needs no introduction for this audience. I was going to call him the father of modern empirical IO, but then John Haltiwanger suggested "godfather" maybe would be an alternative definition. But I want to talk about two papers that he wrote about 20 years ago that are very relevant to this panel.

So the first is the famous BLP paper. That paper was all about taking aggregate data from markets, if you estimate demand, together with assumptions on how firms compete, you can get markups at the firm level or at the product level.

Second, he also wrote a paper with J. Stephen Olley about estimating production functions, and that's the foundation for the supply approach for estimating markups. So regardless of how you estimate markups, Ariel here deserves either the credit or the blame, so...

Second, we have John Haltiwanger. So John is a macroeconomist, and usually when you mention macroeconomics in an IO conference, it's a punchline for a joke, but I think a lot of us don't realize that macroeconomics has undergone a quiet revolution towards requiring empirical evidence and especially empirical evidence for microdata to support theories. That's something that John has done throughout his career and is really a pioneer of doing that, not just with all the empirical papers, but creating the underlying data sets that people now use to try to take macro theories and see if they match the data or not.

Last we have Matt Grennan, so Matt Grennan adds some youth to this panel.

MR. GRENNAN: Much needed. I was wondering why I was here, to keep us awake.

MR. RAVAL: But I just point to you, if you want to know about his work, to the presentation he gave yesterday, which was about looking at price discrimination in hospital markets, looking about how hospitals buy different types of supplies like stints or gloves, and how the markups on those can vary by

bargaining, by search costs, and by other types of frictions.

So the format of this panel is I'm going to ask some questions, and the panelists are going to answer them. So we're going to start with the first question for Ariel.

So I would characterize approaches to markup estimations in three forms. The first is that we might obtain margin data directly from firms, so something the DOJ or FTC could do in a case. The second is you could estimate markups using production data, so that's the De Loecker method that sort of motivated a recent paper saying markups have gone up a lot using Compustat data. And third, through estimating demand, as Ariel did in the BLP paper.

So how would you assess the strengths and weaknesses of these different approaches?

MR. PAKES: Thanks, Davesh. Thanks for having me. I've never been called a godfather before. I guess there's pros -- there's good godfathers and bad godfathers. I hope I'm on the good side.

I actually made, just for this first question, I talked to -- Davesh and I did some email interchange, and I made a few slides just to describe what's going on. So can we move on on the slide? Do

I get a -- is there a clicker for me? That's great. Thank you.

Okay. So I am going to be talking about, you know, how do we measure markups. You know, there is this very important question, which I am not going to be talking on, which is figuring out why the markups are there, okay? You can't tell whether they're too high or too low without knowing how they got there, okay? Why are we having the markups, okay? But I'm not going to talk about that.

So I'm going to give you an example. I'm going to go through demand system stuff, production function stuff, and then what you guys can requisition, okay, what the FTC can requisition. So I'm going to give you an example first.

There was an article in the AER this year, just lately, by Tom Wollman on trucks, okay? So I asked Tom to take down his demand system. His demand system was estimated separately from the pricing equation. He had very good data on demand. He had to do this because he was one of my students, so -- I wouldn't have signed his thesis, but --

MR. HALTIWANGER: Godfather is appropriate, apparently.

MR. PAKES: Maybe that's right, the wrong one,

right?

So take the predicted markup down from the demand system, all right, regress it on supposedly exogenous variables or instruments, okay? And then go to the pricing equation and regress price against the characteristics of the product and wages, which was the only other cost factor, and then look at the coefficient of markup, which should be one if our model is right, okay, and what the R2 is. That's the answers, okay?

It's an amazing fit, okay? If you don't put in time dummies, the R2 is 0.86. The markup is within one standard deviation of one. And if you put in time dummies, it's 0.94, and this is directly from his paper, okay? There's nothing fancy going on. It's just an OLS regression, okay?

So it tells you what I think is true, is that it works very well in the cross-section. So I also asked him to do it over time. So just look at differences in price of the same good over time, okay? So the characteristics are not changing over time, okay? So they're not going into this regression. The only thing that's going into this regression is differences in competition across periods. There are different competitors to these trucks every period.

You still get an R2 of between 0.5 and 0.6, okay, which for a behavioral equation in the social sciences, you know, if somebody is actually setting these prices, is very good. You know, labor has not seen a 0.6 R2 in a long time, okay?

And, you know, I think the reason it works this way, okay, is typically we have really good data on prices, quantities, and characteristics. You know, it's just you know the quantity of cars, you know the prices of those cars, more or less, okay, and you know their characteristics. We don't need input data on cost functions to do this, okay? And, you know, really the open question is the model of pricing, is that good or bad?

Actually, this is trucks, right? So it should be a durable good problem, okay, and it should not be static, but -- and I've seen this time and again. Same one when we did cars, okay? They're a durable good. It should not be a static pricing problem.

But three things happen, which is markups are always smaller, at least every one I've seen in the crowded portion of the market when there's a lot of competing cars with similar characteristics. Markups are higher for high-quality or high-priced goods, and that, you know, just rationalizes the investments in

getting the higher quality, okay? And markups are higher for products where a firm is marketing two products that are competing with each other, just like our theory says.

So it's not exactly right. We know it's not exactly right, they are durable goods, but these kind of arguments make the estimate -- well, make the estimates make some sense, okay?

On the other hand, the problem with this way of doing it, it takes a pretty detailed data set and a lot of time to do it, okay? So you're not going to be able to do it on all of the industries in the economy, okay? It's just not going to be within the feasible set, okay? And I would like to see people like the FTC doing it, but, you know, maybe to make it a little bit easier, I would have suggested having sort of a repository of data on different industries available for people. So that's markups from demand system estimation.

Now I'm going to go to the production side, okay? So this is Jan and his co-authors. So let me take a step back for a second. This literature started with productivity analysis, not with estimating production functions. And if you were honest -- and I did some of this, so I'm not -- you

know, and if you're honest about productivity stuff, you know, really what you're doing is you're getting an index of sales on one side of the equation, and then you're regressing it on either an index of the cost of inputs, not the inputs themselves, or a very loose aggregate of the quantity input, like hours of, you know, very different kinds of labors, high school, university, research, the works, okay? That's what's going on.

And then productivity was just the ratio of the index of outputs over the index of inputs. That's productivity. That matters for a lot of things, but it's not markups, okay? It's an index of sales over an index of inputs. It generates a lot of incentives, okay, for selection and for endogeneity, but it's not productivity. It's not markups. So now let me go to markups.

So how do I get from there to markups? The first thing you need to do is separate price from quantity, okay? And so just using sales is not going to do. And the second thing we're going to need is an elasticity of output with respect to a variable input. So this is the two things that Jan needs, okay?

How do we get these? We estimate a production function and assume it's Hicks neutral technological

change, so the labor or the variable cost coefficient has the same proportional shift as everything else, okay? And then we assume there is an input which is purchased in a competitive market and optimized out in the short run. So you can condition on the quantity of output and the quantity of the other inputs, but conditional on those, this input is going to be optimized out in the short run.

What problems do we get? So the first problem is there isn't a production function for multi-product firms, okay? It's at best a correspondence, right? I have a certain amount of inputs. I can transfer them into different amounts of -- different kinds of output. It just doesn't exist, okay?

So this is often, you know, also true at a plant level, because I've looked at plant-level data, okay? That was surprising to me when I looked at it, okay, but it is true, okay? And even if you did have plant-level data, okay, you know, the firm is not optimizing inputs for the plant. It's optimizing input for the multi-plant firm. And that's a different question with a different answer. Is that clear? So it's problematic, okay?

The other kinds of things that are problematic about it is you really need, you know, the right --

you need to have an index of inputs that's correct in some sense. What is capital, okay? And I need to measure them right. So, you know, you need to aggregate capital stock, and you're averaging things that were bought at very different times, okay, and used for very different things, okay, and aggregate labor stock, okay?

Technological change has to be Hicks neutral, and then, you know, the same problems that arise in productivity analysis, selection and endogeneity have to be dealt with, okay?

Now, there is this huge advantage, okay? So those are the problems, and they're substantial, okay? There's this huge advantage that you can do it for lots of firms, lots of time, okay? It's quick, at least relative to the demand system stuff, okay? If you go to the LRD or whatever other data set you -- I wouldn't go to the Compustat, but LRD I would go to, okay?

And, you know, you can give them the type of data you have. If you believe your, say, materials input is optimized in the short run, you really believe single-product firms is enough, you don't worry about the selection problem, and, you know, there's a reason that there's multi-product firms.

It's not a random draw of firms, okay? If you believe all that, you can do it very quickly. And you can't do that with the demand side. You couldn't do it for the whole -- you just -- it's just not in the cards, okay?

So the third thing is obtaining margin data directly from the firms. So I have less experience with this. The first two I've worked on, okay? This one I've never worked on. The first question is what do you ask them, I think, either for -- and you would ask them different things depending on what you want.

If you were just doing merger analysis sort of conditional on synergies, that's one thing. If you were trying to look at synergies, okay, which they're going to claim when they walk into the FTC or the DOJ, right -- we're not going to do this to raise price, we're going to do this because we have cost synergies, okay -- you're going to be looking at very different things, okay?

If you want marginal costs, you want the inputs -- you need the inputs which contribute to that marginal cost. There's an issue -- there becomes an issue of inputs like marketing that were in the question, because, you know, they do -- they are marginal in some sense, okay, but the returns come

over time, so I need a depreciation rate of something like that to do this, okay?

And when you're evaluating synergies, the arguments -- you're looking more at fixed costs and things like that. So I just looked at the bank data entry, and one of the big things that happens when banks enter is they close branches. When they do mergers, they close branches. That's one of the reasons they're doing it, okay?

So is that -- you know, is that a marginal cost? Is that a savings? It's a savings in cost. At least it's a synergy of some form, okay? On the other hand, it's not very good for consumers sometimes. You might want to take that into consideration. So there's a lot of issues in that, in what to ask for and how to use it.

And then there's also always the question of, you know, when you ask for it, what are they going to tell you, okay? So there's an incentive compatibility problem a la Maskin and Tirole, okay, or Laffont and Tirole. It's a little bit mitigated, I've got to admit, if you ask questions -- you find out margins before the issue that is currently arising is happening with the firm. So if you have emails from prior -- from two years before this when they weren't

thinking of the merger and somebody was telling you something about marginal cost, okay, that would be a different way of looking at it.

So I've told you all the problems, but that's my role as an academic.

MR. RAVAL: So do either of you want to comment?

MR. HALTIWANGER: Sure. It's a little hard to follow the godfather, particularly on this topic, but let me -- one, I'll say I agree with everything Ariel said, and hopefully what I'm about to say isn't just piling on, but I also have some questions, maybe particularly for both panelists, given at least one alternative that I think Ariel didn't mention and has become popular in the recent literature.

So let's talk a little bit more about De Loecker Eeckhout. So just as a reminder, you know, kind of the key equation that they use to identify the markup is very simple, right? The markup is the ratio of the estimated factor elasticity for a variable factor of production, like materials or of labor if you decided labor is a variable factor of production, to the cost share of that factor of revenue.

And the good news is we've got lots of good data on cost shares of revenue, so that's sort of the easy part, but as Ariel has pointed out, we don't actually have, and not off the shelf, we have to work really hard to get that factor elasticity. With all due respect to the De Loecker Eeckhout paper, I don't think they actually have factor elasticity estimates for all the reasons Ariel has talked about, but a key one is that at least that paper does not have the P and the Q data that you need to be able to separate all this out.

I think there's another issue that I think we also -- once you just sort of stare at that formula for a bit, that makes you think, well, wait a second, why do I want to put all the heterogeneity om on the markup side? Why don't I want to put equally as much heterogeneity on the technology side, because in principle, that factor elasticity might actually vary both across firms and time?

So what do we typically do? We end up using the proxy methods, which very much started with Olley/ Pakes, and those methods are somewhat data hungry, and so we often -- to get these, we pool across plants and time, so to get time and variant kind of measures.

So it is kind of piling on. It's all the things that Ariel talked about, but I think on top of that, I think we ought to be thinking about

heterogeneity in technology as much as heterogeneity in markups. And so when I look at that paper, there's something going on, clearly, with the cost shares of variable inputs in terms of revenue, but I don't know whether it's the markup or it's changing technology.

So a somewhat more optimistic view, it is true that there are data sets, even in the United States -- although less in the United States than in other countries -- where we do actually have the P and the Q data. So at least for a limited number of products in the United States -- and this is work I've done with Chad Syverson, but Chad has done it for a whole variety of papers -- there is P and Q data for a limited set of products. And for some remarks I'm going to make later about what's going on in macro, I think we've learned some things out of that kind of work. So there's some hope.

Actually, in other countries, one country I'm working actively with their data is I'm working with Marcela Eslava in Colombia. Nicely, you know, unlike the United States, it's not a Balkanized statistical system, so, indeed, basically the price program is fully integrated with the annual survey of manufacturers. So they actually have detailed P and Q data, not only for outputs, but materials, all

integrated on a regular basis. So it's actually kind of phenomenal, that kind of data. It's kind of economic census every year in that respect.

So again, I think there's hope. There's a variety of European countries that also have this data. So I think we can start going after some of the kinds of issues that Ariel talked about.

And the last thing I wanted to say here is there's a method that's become I'd say increasingly popular, at least in terms of a paper that's getting published in prominent places, and also I'd say the macro literature often uses the estimates of markups or essentially the elasticities of substitution from this literature, and it's really more out of -- it emerged out of the trade literature, and the most recent sets of papers are the papers, for example, the paper by Hottman, Redding and Weinstein, in the QJE.

So that's a paper, just if you're not familiar with it, uses -- you know, what's become increasingly available is transactions-level data, literally UPC code-level data at the P and the Q level. And what do they do? They write down at the product level a pretty simple model basically of demand and supply, but it's at the product level, by the way, so they overcome some of the issues that Ariel was talking

about already.

And then there's huge identification problems, right? Because the question is, I've got -- if I could write down -- by the way, these are a nested CES environment, to do this. They've got to overcome the problem that cost shocks are going to be correlated with the demand shocks. And so what are they going to do? And they don't have the instruments that all these very careful industry studies do.

So they do what Rob Feenstra suggested way back in '94, but perhaps it's more palatable in this data. Why is it more palatable? Because basically they double-difference their equations, and basically they sweep out firm by time effects. And you could -- they argue that they're sweeping out lots of things, and then they say, okay, well the double-difference shocks, they make the assumption -- a pretty strong identifying assumption -- that they're uncorrelated, and that gives them a moment condition, and then they go and they estimate the elasticities.

Now, again, we're often looking for things you can do at scale. They do this at scale, okay? So they take the -- for example, the Kilts data from the Chicago Booth and they've estimated it over 100 markups relatively quickly across a wide variety of

121

product groups, and they've even estimated it at -- some of their estimates are more at the product modular level, so over a thousand markups. So I think that's another horse in the race.

And last issue that I guess I've been struck by is that Hottman, Redding, and Weinstein have argued that this method, in principle -- let's suppose they really did get the elasticities of substitution correctly estimated here. This enables them to extract basically from the demand equation the variation in quality across products, at the product item level.

So, in turn -- and it's also the case that this data has lots of entry and exit of products, so -- and this was one of the -- those who know the Feenstra work, one of the original insights of Feenstra was a way to adjust standard price indices for new product variety.

So why do I bring this up? Because they've developed a method -- and this is more in the more recent papers, the Redding and Weinstein papers -- a method for price indices that adjust for quality change both from product variety and actually for common goods.

And I find it interesting -- and, again, it

122

would be nice to hear I'll say especially from the godfather about this -- is, you know, the methods that Ariel is applying here very much are tied to the hedonic literature in many respects, and that's one way -- and I think a very powerful way -- to adjust for quality.

This offers an alternative. I don't think we fully understand how they compare to each other, but it's kind of interesting that these alternatives -- both these two alternatives yield estimates of markups and estimates of quality change, I think two things we care a lot about.

MR. GRENNAN: It's hard to add. It's a pretty comprehensive set of comments being made here. I guess maybe just two things that are a little bit particular to -- I think I was asked to jump in here because I have a bit of experience trying to apply these method to product markets like medical devices and pharmaceuticals, where the markups tend to be quite large, or at least markups over marginal cost of the actual good, the marginal cost meaning, like, the marginal cost of production and distribution, say, of this actual good tend to be quite large. And I think at least in those types of markets, it tends to be -- you know, estimating this sort of elasticity of an

123

output with respect to an input is likely to be quite small and noisy. So I think it's kind of -- it's a very challenging to imagine that approach applying I think to this set of markets.

In terms of estimating demand and doing it at scale, you know, the paper I presented yesterday is an example where we're trying to do something not quite at the scale of some of the articles that John was citing, but, you know, across quite a wide variety of medical inputs, and there we're doing something that actually conceptually is not entirely different from he was talking about, right?

We're using nested logit instead of nested CES models, right, but hoping that this is kind of -- this plus lots of panel data that allow us for a lot of rich fixed effects is going to capture a lot of kind of first order things that we're interested in in the data. You know, we add -- very particular to our context, we have some instruments to add to it.

But I think, you know, anecdotally, I don't have a quantitative sense of this, but my kind of qualitative impression is that if you aggregate it over a lot of very, very careful IO studies in particular industries, once you have enough data to have very, very fine-grained fixed effects, a lot of

124

times, you know, adding the instruments don't make a huge difference. So there might be something to be said for this type of approach and its scalability.

And I guess I would add to that that, I mean, to me, all of my work is always in -- and most of our work, right -- is in thinking about particular policy questions, right? So this what's underlying the markup, Ariel, you know, decided not to address it, but really that's usually what we'd would want to know.

And so to, you know, to actually address most of the policy questions we're interested in, you know, if you don't have some sort of model of demand, it's very difficult to start to address those questions. So I think that's why you see many of us skewing towards that in our work as well.

MR. RAVAL: Does anyone want to add?

MR. PAKES: Just a couple quickly.

On the linear fixed effects, when you have stuff going in and out, it's not linear anymore. There's a selection problem, and it's problematic. It's correlated with things. So that's true.

You know, this literature on estimating production functions starts with this article by Mundlak about agriculture, you know where the fixed effect was the quality of land, and it didn't change too much over time, but, you know, farms weren't going in and out.

In this case, the products are leaving because they're being obsoleted by better products, and it has an effect on the analysis. On the hedonic stuff, you know, fundamentally, the characteristics base is just an approximation, right? There really isn't a quantity out there. We don't have all the right characteristics to put in the error term.

So if you did it, you know, the other way -- and you could actually do it, you know, we thought you couldn't do it -- so 200 cars with 200 prices for each one is, what, 40,000 cost price elasticity? There's no data set that's ever going to estimate that many cost price elasticities. That's why we went to characteristics. But if you can do it, you know, it's great, but you still need the production function. You still need --

MR. RAVAL: I agree with that.

MR. PAKES: -- and that's where I think most of the problems lie. That's where we don't really have a good grasp on it.

MR. RAVAL: So the second question, and this was brought up I think by both John and Matt, but it's

to John. So I think there's a lot of demand for macroeconomics for an aggregate markup. So the first question is, what role do you think researchers in IO should have in examining this question?

So one way you could think about doing that is just aggregating from estimates from individual industries. So maybe Nate Miller in the audience knows about beer and cement, and Kate Ho knows about insurance companies and hospitals, and Ariel can make his students work on the industries we don't know about, and eventually we get to an aggregate markup.

Another approach would be to do cross-industry studies, which IO has largely abandoned and has been starting to be done by macro and trade people. So what do you think about that?

MR. HALTIWANGER: So there is enormous interest in macro, if you're not aware, in what's going on with the evolution of markups. And I think it's fair to say -- and Ariel has already basically touched upon this -- that for macroeconomists, if we're going to do this, we need to use one of these methods at scale, all right? We're going to need to be able to do this. That's not to say that we shouldn't be learning from the insights from the microeconometrics about both the issues and literally the methods in order to be -- in

order to get there.

So let me try not to -- you know, I'm watching the clock here. I could go on for a while about how macroeconomists are using markups, but let me try to talk a little bit about and use De Loecker Eeckhout actually as both sort of a source, why we're so interested, and then various ways of thinking about things.

So remember De Loecker Eeckhout, the stuff that showed up in the New York Times for De Loecker Eeckhout was about the aggregate markup, right? But actually really what that is, those of you who read the paper, of course, that's the activity-weighted first moment of their distribution of markups. And actually their paper is very much as much about what the "aggregate" markup, the first moment is doing, as they -- they have quite interesting things to say about the evolution of the distribution, changing dispersion, skewness, the connection between changing skewness, and the first moment, and so on.

So, again, I think -- and I'll say just enormous interest in that, and, indeed, you know, where this has sort of led, there has become great interest, partly from De Loecker Eeckhout, that if we took -- the one possibility is that competition has

become more imperfect in the United States over time. We are a less competitive economy, which is, you know, a controversial statement, and this is related to the question about, where are these markups from and what might be driving them?

Now, there's a parallel literature that's emerged in macroeconomics that I'm almost hesitant to bring up in this audience, because actually there was a Jackson Hole conference just recently very much about this, and it was very much focused on changing concentration, and there are lots of macroeconomists who have been using industry-level concentration measures to shed some light on this.

Two of the people at that conference were from the IO community, particularly -- and probably more than that, maybe I'm forgetting others -- but Chad Syverson was there and so was Carl Shapiro, and both of them -- I think Chad called using kind of concentration metrics and said -- one of them called it the original sin and the other one called it the forbidden regression. I don't know which one.

So both of them -- so basically they said you have to be incredibly careful about using this outcome variable and they came through examples. You don't know which direction -- even as Chad walked through

129

examples, concentration may go up or down with the changes in competition depending upon the model and the structure. So I mention that not to advocate for the use of concentration, particularly in this audience, but rather just to indicate the enormous interest.

I wanted to build on a little bit back again to De Loecker Eeckhout, that they're really much more than about first moments, they're about the distribution of markups, and I wanted to talk briefly about why macroeconomists are so interested in the evolution of the distribution of markups, what they are, how they might vary across time, countries, industries, and the like.

And also I think we've learned something from the literature -- I'm just about to talk about it -- that provides indirect evidence about what might be going on with markups. So what literature am I referring to? So one of the areas that's become a focus I'd say in the last -- particularly the last decade, although it's an older topic than that -- is misallocation. So what's the -- it's become kind of a working hypothesis that especially if we're trying to explain differences in economic performance across countries, but often also within countries over time,

130

that deteriorations in aggregate productivity reflect changes in misallocation.

And the paper that's probably gotten the most attention, certainly the most cites, I think, is a very nice paper by Hsieh Klenow, and I want to talk very briefly about the Hsieh Klenow paper, both because of its insights, but also then I'm going to come back and talk to you about what might be going on in the data that we've been looking at that actually might be driven exactly by the De Loecker Eeckhout changes in markups.

So here's sort of the Hsieh Klenow 101 for those of you who are not familiar, really quickly. So like most macroeconomists, they write down a very simple model of the production technology and the demand structure. So critically they use CES preferences and end up actually, even though they -- on the production side, by the way, they allow for heterogeneity in production elasticities, which they measure from cost shares, by the way, like growth accounting. On the demand side, they take one number. The elasticity of substitution is four, okay, and so the markup is 1.33 in their very simple calibration.

So here's the key insight, by the way, and it's a very powerful insight -- and in some ways too

131

powerful as I'm going to argue in just a second -- but nevertheless, I think it's very insightful, including for this discussion about the evolution of markups.

So if you write down a -- write down this model and you think about the -- what we'll call the frictionless benchmark. What's the frictionless benchmark? Where marginal revenue products are all equalized. Well, you could say -- we've actually known about this for a long time, but they really emphasize this. If marginal revenue products are equalized in exactly this setting, but inside industries, measures of revenue dispersion, exactly the measures that Ariel was just talking about, the sales per unit input should exhibit no dispersion. It's a consequence of margin revenue products being equalized.

But as Ariel and I know -- we've been doing this for a while -- there's enormous dispersion across businesses in these measures, and not only that, in the classic Olley/Pakes paper, it wasn't just that they did great things in terms of estimating the production or sales function, but they showed that as the telecommunications equipment industry underwent changes in the economic environment, there were important changes in measures of allocated

132

efficiencies. So these measures are quite indicative. So the question is, how do we reconcile the Hsieh Klenow view with I'll say maybe the Olley/Pakes view of TFPR dispersion?

So, remember, let me just go -- I didn't quite finish the punchline of Hsieh Klenow. So what did they do? They say, well, look, we see enormous dispersion in revenue productivity dispersion across firms and plants in the same industry. It must be driven by wedges, some sort of distortion.

And so they found, for example, that revenue productivity dispersion is much larger in China and India than the United States, and they -- given their strong assumptions, they could literally back out the distribution of wedges, and then they could actually do a calculation that said, here's all the allocated inefficiency in China and India that resulted.

Now, as we thought further about this, we realized there's a whole host of things that might be driving revenue productivity dispersion above and beyond the kind of wedges and distortions that they're talking about. Some of them are things like adjustment costs, and you say, well, gee, how do I distinguish between an adjustment cost and a distortion? Well, lots of us have been certainly writing down models, dynamic models where even a social planner faces a certain amount of adjustment frictions for labor or adjusting the scale, and I know there's lots of interesting things in the IO literature about this. So one tension in the literature is how to back out all the -- I'll call it the wedges and frictions that are part of the environment versus the residual wedges that are present.

So now let me get back to the distribution of markups. So I mentioned that there are at least some products in the United States for which we have the P and the Q data, and there's lots of data sets around the world, I mentioned Colombia, but there's data sets in Europe and so on. And not that we've solved of Ariel's problems, but -- when we do this, but when we do this, we can't -- we have a lot at least at being to estimate the production technology, all right?

And so because why? Because we can -- we can compute a measure of Q. We still have multi-plant firm -- multi-plant -- excuse me, multi-product plant issues to confront, and actually, recently in my work with Marcela Eslava, we've been going after just that as well, but I don't want to go down that path right now.

But here's one of the things that we have found, and this is particularly in my work with Chad Syverson and Lucia Foster. So one is the measure of TFPR, and I've already told you that TFPR, in principle, may, under the Hsieh Klenow assumptions, only reflect wedges. But we've found when we've used the P and the Q data, that the underlying -- I'm going to call it TFPQ, the underlying productivity -- and by the way, that measure really may be more -- a better way to think of it is a competent measure of both productivity and demand or quality, but let's just call it TFPQ for right now.

We found an incredibly high correlation between TFPQ and TFPR. I think that correlation is really important because it actually helps explain why Olley/Pakes found some of the findings they did; that, indeed, the TFPR measure that they were essentially using is highly correlated with the TFPQ measure, and that's why the Olley/Pakes -- amongst this group, probably the proxy method is the most famous, but their very simple and very nice decompositional allocative efficiency was also a very powerful result in that paper, and I think it's partly being driven by the fact that we find such high correlation between TFPQ and TFPR.

So the question is, why is there such a high correlation between TFPQ and TFPR? And I am going to bring this to a close and bring it back to markups. So one possibility, if I -- is that it's correlated distortions, right? I've got wedges out there, and there's some black box reason. That might be. I think there's probably some of that, but I actually think that we have much better explanations of this correlation.

So one of them is actually something I've already mentioned, is adjustment costs. As soon as you write down an adjustment cost model, then it's going to be the case that as a firm gets hit by a shock, it's not going to adjust completely, it's going to take time, and as a result, even in the CES framework, you know, what's driving this Hsieh Klenow result is there's actually a negative unit elasticity between P and TFPQ, and that's going to disappear in an adjustment cost model.

But what's another powerful explanation that I think actually may be playing a huge role is variable markups and, indeed, markups that increase with TFPQ. Do we think there's evidence of that? Yeah, actually, there's a very nice paper -- it's actually more of a theory paper, but I was really struck by its

discussion on the evidence, this paper by Dhingra and Morrow coming out in the JPE about variable elasticity models. And it's actually much more about how to think about entry/exit models with variable elasticity.

But it partly, in its motivating section, it talked a lot about what they regarded as the indirect evidence that markups are variable across firms, and actually tend to be increasing in size and fundamentals. And basically what they cited over and over again was the incomplete pass-through literature. There's lots of literature that suggests that businesses, when they get hit by costs, do not pass those costs on completely.

One good way of explaining that is, indeed, variable markups. So now I've come full circle. So I think macroeconomics cares a lot about, not surprisingly, the misallocation. To be able to measure misallocation, actually we need to understand the shock process that's hitting businesses. We need to understand the heterogeneity and the technology. We need to understand the heterogeneity and markups to be able to get at all these kinds of things.

And so guidance from the IO literature -- so back again to your original question -- are the

137

macroeconomists hungry for good estimates of the distribution of markups and how to think about this? You betcha. And not only that, tell us about the distribution of markups, but also how they vary with key fundamentals like TFPQ. And I'll stop.

MR. RAVAL: Do either of you want to comment?

MR. PAKES: Just one quick comment on the Olley/Pakes stuff.

So it may well be true what you've said, that correlation with TFPQ, but there, if you look at Olley/Pakes, look at the mean -- the average productivity as you go along. The distributional stuff that John said is correct, but the mean actually doesn't increase, and this is probably the fastest moving -- it's telecommunication equipment. Stuff was moving very fast. There was a lot of technological change.

And I think the reason for it is price was going down, because they opened up the market and there was competition. So, I mean, it's got to -- it depends on the question you're asking, what was going on, okay? So just that.

MR. RAVAL: All right. So the next question is for Matt Grennan.

So some of the industries we study, such as

138

pharma or high-tech, are characterized by a high fixed cost and low marginal cost. So is a markup over marginal cost even relevant for these industries, first of all? And second of all, you know, one way to view the De Loecker Eeckhout evidence is maybe the relevant points of fixed cost versus marginal cost is changing. And so if you start moving towards a more high-fixed cost technology, that could lead to increases in measured markups.

So do you think that's what's going on, and what should we be doing about it?

MR. GRENNAN: Yeah. So I think this question of, you know, the role of fixed costs, so I think -- you know, at least when I hear that, I think, you know, costs of research and development or maybe costs of adopting new, expensive technologies, or at least kind of quasi fixed costs, like kind of sales, marketing or management, you know, and are those leading to some rise in markups.

I mean, I'm not sure that we have great evidence. I think it's a very interesting hypothesis that we should probably all be exploring to some degree, but, you know, in terms of -- so I'm not sure I can answer, you know, are those at the root of things, but I do think that keeping those separate

139

from markups relative to the marginal cost of production and distribution of a good is important, right, because they're very conceptually distinct in terms of, you know, what they tell us about implications for short- and long-run efficiency and responses to various changes in policy or in the economy.

I'll give -- maybe two examples might be helpful, right? So in hospital markets or medical markets, right, having a good sales force for, say, selling a pharmaceutical or a medical device seems to be an important thing for selling a lot of the product, right, for generating sales in those markets, and probably part of this is that there's some value-added service component sometimes.

In medical devices, it's part of just how distribution works. You know, there might be part informative aspects to this, right? You're getting the word out about these technologist and how they're best used and so on. And, you know, there's likely in many cases also persuasion in these activities, right?

And if you have, you know, an oligopoly or a monopoly industry, as we often do in some of these, you know, there's likely maybe some business-stealing aspects to those. And so in that case, it may be very

140

inefficient, this spending, right, and maybe even if you're persuading people to allocate things to the wrong patient during the wrong circumstances, perhaps even value-destroying, right? So, you know, I think you'd want to keep those things separate when you're thinking about markups in that case.

The other might be, you know, in these industries, as in a lot of intermediate good industries, the prices that you're looking at are often negotiated, right? So this markup not only has something to do with, say, demand elasticities, competition, and marginal costs, but also where some bargaining parameters tell you that the price is ending up in between some bounds that are created by those other forces, right?

And if you think about -- you know, look at any of the estimates that we've been getting from these sort of models, there tends to be a lot of variation left in this bargaining residual that's kind of explaining where prices are ending up, right? And kind of qualitative evidence in my experience suggests that, you know, what's driving these? A lot of things like managerial skill, effort, maybe information, and, you know, if this is just transfers, then investment in this kind of bargaining, you know, effort is just pure social waste, right?

If it's not transfers, then this affects allocations either in the immediate market, because you can't contract on quantity, or in downstream markets because this ends up being an input cost for those downstream markets. Then, you know, investment in negotiating a better price by suppliers is unambiguously I think a bad thing, typically. I mean, there's caveats for oligopolies, and if we're aligning prices with marginal costs and so on, but probably likely a bad thing.

But on the buyer side, investment in negotiating better price is probably also conversely unambiguously good, because you're driving prices closer to cost potentially and increasing at least short-run allocative efficiency in that sense.

So I think that, you know, the big takeaway to sort of answer the question, yes, I think the markups -- the traditional markups are still very valuable in these cases, and I think that these fixed or quasi fixed costs, you know, we should be thinking about them, but we should be thinking about them I think as distinct components in how they interact with these markups.

MR. RAVAL: Do either of you want to comment?

MR. PAKES: I actually -- my only comment -- maybe two things. We're not -- there's a sense in which we're not thinking about dynamics, and, you know, in certain industries, that's where -- if you were doing a merger in pharmaceuticals, that's the first thing I would worry about. I would worry about the R&D policy of the industries.

I thought of that mostly because of what you said about the bargaining thing. You know, in the hospital thing, which you've worked on, the thing that I think is most interesting about the bargaining thing is it splits the profits, and it's going to determine investment incentives. Depending on where it splits the profits, we're going to see, you know, arms races or we are going to see savings in costs, and we're not focused on that, and I think we should be focused on that. I'll just leave it there.

And, you know, the reason we're not focused is it's difficult, but, you know, we can start with something, like reduced-form stuff, anything, to get a handle on what's really going on with the investment stuff, so...

MR. RAVAL: So the next question is for John.

So, as many of you know, the FTC has 6(b) authority to subpoena firms, and we used to ask firms

for data on profits, revenue, and other variables by line of business, but we haven't done that since the 1970s. So is there data that companies have that is not already collected by the Census or someone else that would be useful for markup estimation, for understanding competitive conditions? And if so, what would be useful and should we try to do that?

MR. HALTIWANGER: So a really good question. By the way, it's very difficult to say we don't need more data, but let me talk about I think where we stand relative to when the line of business data were collected.

So actually I went back and looked at a very nice paper written and published in 19 -- I think it was 1991, by Ravenscraft and Wagner, and it talked about the value of the line of business data, and actually it compared it to at that time, what was considered a new entrant in the market, the LRD, but it also compared it to Compustat and so on.

So here's the good news. Let me try to give you a little bit of a sense of I'll say the enormous progress the Census Bureau has made, in particular, in building business-level data sets. So the LRD, the one that, you know, Ariel and I started working with a very long time ago is manufacturing only. It's built

on the Annual Survey of Manufacturers and the Census of Manufacturers.

It had reasonably good establishment-level identifiers, and kind of okay firm-level identifiers. So you could do some analysis, but we didn't really have, even though underlying this there was a Census Business Register, we really had not built a longitudinal version of the Census Business Register. That now exists. It's called the LBD, and you could say we're not very creative about coming up with new acronyms. It's the Longitudinal Business Database.

It's a remarkable database. It's from administrative data, and it tracks -- and survey data, I should say, because -- you'll see why in a second. It actually tracks every establishment in the private sector over time, and it has all the parent linkages. And where is that all coming from? The parent linkages are coming from the economic censuses and the Company Organization Survey.

And so you can do a remarkable amount about firm versus establishment dynamics. Now, currently, that data set -- this is going to eventually get back to your question -- that data set, one of the primary variables -- I don't know, we've got incredibly good location data, we've got the organizational structure

145

data that we just talked about. In terms of sort of the economic outcome variables, the key variables that are available are employment, payroll, and revenue. The employment and payroll data come from payroll tax reports, and the revenue data comes from business tax returns.

Now that data, it turns out, Census is getting the complete dump of all forms of business tax returns. So there's lots of information on costs of materials, actually other kinds of costs. There's even -- you can certainly -- you can build an accounting profits measure from the administrative law. So in that sense, the data are sitting there, and you need folks to kind of come and invest and spend time building that up.

And you say, well, how hard can that be? So the LBD is something that was created in the last decade or so, and then just a few years ago, you know, one of my research assistants and now co-authors, Robert Kulick, we had him add the revenue data. It took him three years to add the revenue data, to be able to sort through -- he had to understand all the different tax forms and the fact that Census was changing the way they were doing things over time and so on.

146

But anyway, the point here is I think we are in much better shape than we were back in 1991 when Ravenscraft and Wagner wrote their paper on I'm going to call it the core accounting profits notions, the notion that there's -- and to be able to do things at both the establishment and the firm level. And I'll say as well, because of that -- and there has not been enough of this done as well -- but an enormous amount can be done on studying merger activity and changes in ownership structure. The data are there. People were beginning to push the data hard in that direction.

Now, what do I think is really missing? And my question is, you know -- so I don't think we are missing the kind of line of business notion is the point. I think we're doing a pretty good job, and I'm going to go ahead and emphasize it. The industry codes are fantastic, okay? They're state of the art industry codes. Why? Because they come out of the economic censuses, where you really ask the detailed questions. So that enables you to track business activity.

So what are we missing? Well, we could do a lot better on capital if we possibly could. So capital is a really hard one. Ariel sort of talked about this. So lots of, you know, basically measures

147

of -- there are some measures in the accounting data on capital expenditures, and there's book values and so on, but it's pretty crude. So helping just figure out what's going on with capital, it would be a big deal.

The other one we've already hinted at, what do I think we're really missing in the United States, and the question is whether, you know, there are gains from trade here somehow or another, is P and Q data. That's where we're really in deep water. The set of products at Census at least for which you can do P and Q is, you know, I think we -- we've done interesting studies, but I was talking earlier about going to scale to be able to look at various things. Can you go to scale -- go at scale? No, it's only 150 limited products in Census where you have the ability -- and it's only every five years anyway -- to do P and Q.

Another place we don't have much good data -- and I don't have any idea whether your data would -- data you could get ahold of or what folks in this room work with -- but we know very little about the supply chain. So we don't know who buys from whom, and so even some of the recent work that, for example, Chad has done with Olley is work that was in -- you know, they started getting indirect things in terms of

148

vertical integration about who -- when, indeed, they saw evidence of firms integrating, they tried to back out who was buying from whom based upon physical location information, so sort of indirectly.

So on the one hand we've made enormous progress, I'd say, in now really truly comprehensive data sets, tracking all firms and establishments in the United States, and I think it's actually still underexploited. Lots of papers have been written, but I think an enormous amount of things could be done.

The big missing pieces are P and Q and supply chain, and then, more generally -- and, again, here's where we could again make a -- it would be great if we could somehow make progress on this. So if you're not already aware, the United States, BLS, Census and BEA, the three primary agencies that put together the key national indicators, like GDP, they can't share their microdata, and so there's great data sitting off at BLS, everything from occupational data to price data for that matter, that could be integrated, in principle, and you can't do it right now under the current legal environment.

So while it would be great to think about partnerships and so on with the FTC, from my -- I'll say from my vantage point -- and maybe even yours --

if you could get at integrating the BLS, Census and BEA, and I just want to mention, BEA is also sitting on top of fantastic data on FDI and multinational activity, and, again, that data can't be integrated. So it's kind of crazy that we are almost unique among the advanced economies where we're so Balkanized and you can't bring all the pieces of the data together.

MR. PAKES: Can I say one thing?

MR. RAVAL: Sure.

MR. PAKES: So when we were doing the LRD, John -- by the way, we all owe John a great deal of accolades, because he's one of the guys who has really gotten the data together in this country. I was there at the very beginning and then did other things, and John just kept doing it.

MR. HALTIWANGER: Some minor things like BLP, right?

MR. PAKES: But when we were at the Census and we were talking about setting up the regional data centers, we had people who could go in and do this stuff, there was a lawyer there, and he said, you know, if one of these numbers gets out in a court case, in a merger case or something like that, the firm can shut down the whole Census, because -- and that's the reason that they're so worried about it.

I mean, I was really amazed when they let us do the LRD after all that. We got this lecture on what could happen if something got out, and they still allowed the LRD, but it is a serious issue. I mean, I would be all for -- I mean, if you could get the price data at the BLS together, the BLS has very good price data. I've worked on the Consumer Price Index before. It's like they actually sample actual goods and all their characteristics, and then they go back and sample the same good again to find out what happened to its price. That's how you get a price index.

But I don't know how, you know, you can get them to merge it. The lawyers won't let you, I don't think.

MR. GRENNAN: I mean, just one thought, harking back to this issue of thinking about investment incentives and all these different pieces of data. I think one of the things that keeps us from doing, you know, more work I think on the investment in these kind of fixed or quasi fixed costs is not only -- you know, it's hard on the kind of conceptual side and theory side, but also difficult on the data side to map these -- whatever data you might be able to get on these areas into the product markets that you think they're targeted at.

And so I think that's another -- to the extent one is asking for data or trying to construct data sets, you know, that I think is another thing that is very difficult and kind of connecting some of the upstream costs or the kind of costs that aren't necessarily well allocated to product markets that they're targeted at, I think that would be a very useful thing to have.

MR. PAKES: It's also very important for vertical. When you guys are doing vertical integration, that's one of the issues. The issue is how the upstream guys' investment incentives correspond to the downstream guys.

MR. RAVAL: So we've got maybe five minutes left, so do any of you have any concluding remarks, something you wanted to say that has not been touched?

MR. PAKES: I have one thing. I talk too much. I have one thing, which is I really think rather than focus -- I mean, I understand the focus of FTC and DOJ on markups, and for short run, for things like mergers, perhaps, but, you know, I think the real issue is what's underlying the markups. I mean, Matt said this, but, you know -- yeah, that's really the question.

The question isn't -- you can't answer the question of whether, you know, maybe we'd increase

markups, but we don't know whether that's good or bad. It may well be -- you don't want to not have Google, right? You might want to do things about it, okay? You don't want to not have the drug companies. You might want to change the rules somehow, but you don't want to not have the drug companies.

And in order to understand either the tech sector or the biotech sector, I don't think it's possible to understand it without knowing more about dynamics, and we're not doing that.

MR. HALTIWANGER: I'm fine.

MR. RAVAL: All right. So I had come up with eight questions, and I could only ask half of them, so we could probably continue this panel on for another hour or two, but we are just out of time, and so we will conclude.

(Applause.)

MR. ROSENBAUM: I'll just give a quick thank you to everyone, our scientific committee, all the panelists, moderators, discussants, presenters, thank you very much, and hopefully we will see you next year at the Twelfth Annual Conference. Thank you.

(Applause.)

(Whereupon, at 12:41 p.m., the conference was concluded.)

CERTIFICATE OF REPORTER

I, Jen Razzino, do hereby certify that the foregoing proceedings were digitally recorded by me 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; that I am not a relative or employee of any attorney or counsel employed by the parties hereto, not financially or otherwise interested in the outcome in the action.

JEN RAZZINO, CER
Court Reporter

A

P