FEDERAL TRADE COMMISSION

THE TWELFTH ANNUAL
FEDERAL TRADE COMMISSION
MICROECONOMICS CONFERENCE

DAY ONE

Thursday, November 14, 2019
8:45 a.m.

Federal Trade Commission
Washington, D.C.

FEDERAL TRADE COMMISSION

I N D E X

P R O C E E D I N G S

- - - - -

MR. KOBAYASHI: All right, good morning. Because of my upbringing, I feel compelled to start on time. My name's Bruce Kobayashi, and I am the Director of the FTC's Bureau of Economics. I'd like to welcome you all to the 12th Annual FTC Microeconomics Conference.

For those of you who are from outside the FTC, I want to say a few words about our agency and the Bureau. As you probably know, the FTC is an independent agency, and it has two primary enforcement missions. One is consumer protection. The other one is competition. We're attempting to prevent business practices and conduct that are anticompetitive or deceptive and/or unfair to consumers.

The FTC also has a broader mission to enhance informed consumer choice and public understanding of the competitive process, and one way we do that is, of course, we produce research and reports, which is also a big part of, I think, the BE mission and I think, you know, from my review one of the more important parts of the BE mission.

The Bureau of Economics is about a little over a hundred people, including about 80 Ph.D. economists.

That makes us one of the larger groups of microeconomic economists in the Federal Government, and we do a lot for the agency. We support directly both the antitrust and consumer enforcement missions; we provide economic analysis in support of investigations and litigation; and we apply in many cases cutting-edge economic analyses, both theoretical and empirical, to these cases.

BE also supports the FTC enforcement mission as well as the mission to sort of inform consumers by producing and publishing a lot of high-quality, cutting-edge research that, in effect, is to be applied to our sort of direct support of enforcement.

Today's conference, like its predecessors, complements our robust economic research program. The conference features cutting-edge academic research with extended discussions of their relevance to real-world economic applied problems. The conference organizers and the scientific committee has again put together an outstanding program, a great set of presenters, discussants, and panelists, and I'm really looking forward to the two days.

So before the first panel, I want to do a couple of things. One I have to do, I have to do a bunch of announcements, and I'll do that last. But I

want to sort of do an extended thank you to all the people involved in sort of making this happen. It was a lot of work, and I'm not going to be able to mention everybody, but I want to mention a bunch of people by name.

First, I want to thank our cosponsor, the Tobin Center at Yale University and its faculty director, Steven Berry. I think they're new this year, and it's great to have that support.

I want to thank all of the BE staff who helped with planning and organizing the conference. In particular, I'd like to thank the conference organizers, Ted and Antara. I would like to thank Alex Avramov for outstanding work once again for putting together the conference logistics and sending me timely emails so I know what to do.

Tom Koch, Miriam Larson-Koester, James Thomas and Will Violette for assistance with the scientific committee in selecting papers and helping with the sessions, and, you know, all of the BE staff economists who spent time in giving feedback on the many submissions that we got for this conference.

I want to especially thank our scientific committee, Panle Jia Barwick from Cornell, Joel Sobel from UC San Diego, Mark Schankerman from the LSE for

serving, and I also want to thank our wonderful BE administrative team who always do incredible work behind the scenes to ensure that this thing goes off well. That includes Maria Villaflor, Kevin Richardson, Neal Reed, Constance Herasingh, Priscilla Thompson, and Tammy John. They really sort of do a lot of work in getting everything set up.

I want to thank our research analysts and statisticians for helping with registration, and always the FTC media team, the Office of Public Affairs, the Office of Executive Director, especially this year for fighting with the OCC about conference space here, and the event planning staff. The FTC is a small agency, but these things are -- I mean, our technology is great, and we are blessed with great support.

All right. So the last thing I have to do is I have to read a bunch of things. Don't be alarmed. I've done a lot of these and have never had to use these. So there's some administrative details.

First, please silence any mobile phones or other electronic devices. If you must use them during the workshop, please be respectful of the speakers and your fellow audience members.

Please be aware that if you leave the building, the Constitution Center, for any reason during the

workshop, you are going to have to go back through security, so that should deter you from leaving. Please bear in mind that you'll have to do this and plan ahead, especially if you're on a panel after lunch, so we can do our best to remain on schedule.

Most of you received a lanyard and a plastic badge. The value is pretty low, but we do reuse them, so please return them at the security desk when you leave.

If an emergency occurs -- this is the scary part -- that requires you to leave the conference center but remain in the building, follow the instructions provided over the building PA system. If an emergency occurs that requires the evacuation of the building, an alarm will sound. Everybody should leave the building in an orderly manner through the 7th Street main exit. After leaving the building, turn left and proceed down 7th Street and across E Street to the FTC emergency assembly area. Just follow Ted. Remain in the assembly area until instructed to return to the building.

If you notice suspicious activity, please alert building security. Please be advised that this event may be photographed 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 there is no right to be forgotten here. If you want that, go to Europe or Argentina.

Restrooms are located in the hallway just outside the conference room. I'm almost done. There is a cafeteria here. It's actually decent. Breakfast, 7:30 to 10:00. Limited menu 10:00 to 11:00. It reopens for lunch at 11:30 to 2:00. So, actually, it's decent, but I think we have lunch here, so it's irrelevant.

Now I am going to turn the podium over to James, who will start the first panel. Thank you, and I look forward to a great conference.

(Applause.)

MR. THOMAS: Thanks, Bruce.

So my name is James Thomas, and I'm an economist at the FTC. We'll get started with a paper session organized by Joel Sobel, which includes two papers related to tacit collusion. So we are going to have 25 minutes for each paper, ten minutes for each discussant, and then ten minutes for audience questions.

So, first, we have Vincenzo Denicolo from the

University of Bologna presenting a paper titled "Artificial intelligence, algorithmic pricing, and collusion." Thanks.

MR. DENICOLO: Good morning. First of all, let me start by thanking the organizers for having us. This paper is joint with Emilio Calvano and Sergio Pastorello, who are colleagues of mine at the University of Bologna, and Giacomo Calzolari, who is now at the European University Institute in Florence.

As the title says, it is about algorithmic pricing. I believe everybody is familiar with the notion that algorithmic pricing is becoming more and more prevalent. If you book a flight or you buy something on Amazon or, increasingly, if you go to a gas station, chances are is the price that you see has not been set by a human decision-maker but by an algorithm.

So in the light of this, several people have raised concerns that these algorithms may collude even if they have not been designed to collude, even if they have not been instructed to collude, because to the extent that collusion is built into the algorithm by design, I would say that that would create not really new problem compared to collusion among humans.

New problems may, however, arise in the

possibility that we have really tacit collusion among algorithms when the algorithms learn to collude autonomously. It is without having been designed in order to collude. So that is the concern that we address in this paper, and the question, in particular, that we ask is, how real is the risk that these algorithms may autonomously learn to collude?

And, of course, that is a question which has important policy implications, and depending on the answer that you give to these questions, then you would go for a different kind of policy. So basically if, as some people do, you believe that we don't have any real problem here, that algorithmic collusion is something we say is really unlikely, then you would go for a laissez-faire policy.

If you instead believe that algorithmic collusion is very common, very easy to achieve, then you might want -- as some people have suggested, you might want to regulate this particular sector and have the firms use only those algorithms that have passed kind of ex ante examination like we have for drugs, for example, that can be brought to the market only after being approved, having been approved by the FDA. And in between, somewhere in between these two extremes, there could be scope, of course, for antitrust policy.

So in particular, I mean, algorithmic pricing, as I said, has become more prevalent, but it is not new, okay? Airlines and hotels have been using algorithms for decades by now. What is new, however, and what, you know, makes the problem of autonomous algorithmic collusion nontrivial, is the software, okay? The software is no longer -- or in addition to the rules-based software that we used to have in the past, it is no longer rules-based, but it is based on artificial intelligence, and it is based on reinforcement learning.

For those of you who play chess, it's pretty much like the comparison between Stockfish. Stockfish is a standard program, it says. Every serious chess player used Stockfish or the equivalent of Stockfish to make analogies, you know? Now, Stockfish is built in this way. It is -- built into the software, there is a function that allows the program to assess any possible position that may arise in the course of playing the game, okay? So that is to be fed into the program.

The new software program -- chess programs, which are based on artificial intelligence, such as AlphaZero, which was created by a team of scientists at Google and which beat Stockfish -- by the way, Stockfish used to be, you know, the standard. So when

was the last championship among humans, okay? Commentators, when they wanted to assess "objectively" the position around Stockfish, okay? And when they wanted to praise the human players, they would say, well, they play like Stockfish, okay?

But Stockfish was beaten easily by AlphaZero. In a round of 100 matches, AlphaZero won about a quarter and drew all the other ones, okay? Stockfish never managed to beat AlphaZero, and AlphaZero is based on artificial intelligence. It is not built with a function that assesses the position. It learns everything from scratch, okay? And it is trained in self-play mode or against other programs until it learns how to play, and it does it very effectively. Now we have a software algorithm that does the pricing in pretty much the same way, okay?

Now, before showing you what we do in the paper, this is a brief summary of the early debate. In the early debate, some people, as I said, were mainly unconcerned. They believed that we don't have really anything really new here. Other people instead jumped to the conclusion that algorithmic collusion is inevitable basically on the basis of the argument that algorithms can change the price much more frequently.

Now, as the frequency of interaction increases,

of course, then the gains that you may obtain by deviating before being punished in case of tacit collusion becomes pretty small, okay? The rhetoric is firms may react before consumers do, okay, and, therefore, there is no gain at all from defection. Actually, that is not possible literally, but the meaning is that the length of the period becomes really short.

However, the skeptics would argue that the size of the discount function, because the length of the period magically boils down to the size of the discount factor, is not the critical factor in limiting the possibility of collusion. Much more important is the ability of humans or algorithms to solve a complex coordination problem under conditions that may be changing or with -- in the absence of explicit communication among them.

So how can we approach this question, so the question of assessing how real is the risk of autonomous collusion among algorithms? Well, to me, the most natural -- I meant applied series. So the most natural approach would be a theoretical approach, okay? So let's study how these algorithms work.

Unfortunately, this approach doesn't work. It's not feasible, basically. We explain in the paper

in some detail what problems arise with this theoretical approach, what progress has been made, and how far we are from being able to say something useful for policy by following this theoretical approach.

Another approach could be empirical, but that, too, is very hard. My discussant today will be able to tell you better than me what problems there are in detecting collusion by looking simply at the market data, and, in particular, one of the of problems that we may face when trying to detect empirically algorithmic collusion is that firms do not disclose the type of algorithms that they use, okay? So that is something that must be inferred from the data, and it is quite hard.

So what we do is we call it experimental approach, but it is actually numerical simulation, okay? So we build reinforcement learning algorithms, we let them interact repeatedly in an artificial market, and we see what they do, okay? Of course, this kind of experimental approach has problems of external validity. What we find in this experiment may not apply to the real world. There may be two basic issues.

One problem could be that the environments that we use in our artificial experiments may not be

representative of the real world markets, and the algorithms that we use may not be the same as firms do in the real world, and we will discuss how we deal with this problem in the presentation, but before doing that, let me jump to the findings that we arrive at in this paper.

Basically, we find that even simple reinforcement learning algorithms -- and we examine algorithms of the Q-learning type and I will explain in a moment what they are -- learn to collude systematically, okay? Collusion is not perfect. They basically learn to collude partially, meaning that they do not coordinate, or not always, on the monopoly price. They coordinate typically on lower prices, but they do it quite systematically.

And these high prices that the algorithms eventually learn to charge are rational, meaning that, as we know, starting from high prices, there would be an incentive to undercut, okay, but that incentive would be countered by punishment, and I will show you what the punishment strategies look like, okay?

What is striking is that the algorithms learn not only to charge high prices but learn to punish in the way that I will show you from scratch, okay? They do not know anything at the beginning of the

iterations. All they know is that they have to maximize profit, okay?

So they know the target. They are told what they can condition their prices on, okay? That is, we allow, of course, the algorithms to condition the current prices on the past prices -- otherwise, we couldn't have collusion, okay? -- but that's all tied to the algorithms. All the rest they learn from scratch, okay, and they learn to collude without communicating with each other.

Let me skip this in the interest of time. As I said, we focus on Q-learning algorithms, and it may be worthwhile to spend a few minutes telling you -- explaining what the Q-learning algorithm does. So basically the engine of a Q-learning algorithm is a table -- a matrix, if you like -- like that. On the rows of this table, you have the current state of the system, okay?

So this example is an example in which the algorithm has a one-period memory, so it can recollect the prices charged by itself and by the other algorithm -- that is an example of two competing algorithms -- in the previous period, okay?

Now, for each pair of prices in the current period, the algorithm has a number of functions, okay?

Now, we have a table here instead of a function if we discretized the prices. Q-algorithms require the action space and the state space to be discretized, okay? So one thing that we do is we replace a continuum as the action set with a discrete set of prices, and so we have a discrete set of possible states.

For example, in our baseline specification, we have 15 feasible prices, okay? So there would be 15 columns in this matrix, okay, and 225 rows in the matrix that is pairs of prices that may have been charged in the last period.

Now, for each of these entries in this table, the algorithm has a number which tells him how valuable it is charging that price, 10 in this example, given the past prices, okay? Now, how is that number determined? The idea is that you start from arbitrary numbers, okay? In the baseline specifications that we use, these numbers are actually the discounted profit that the algorithms would make by using -- by charging a certain price on the assumption that the other algorithms would randomize uniformly across all possible prices, because we start exactly by letting the algorithms do that, but basically, at the initialization of the matter, it is arbitrary, and we

do -- I mean, we do a robustness analysis. If we allow for enough experimentation, as I will tell you in a moment, the way the matrix is initialized is not that important, okay?

And then starting from these arbitrary values, any time the algorithm visit a cell, if there's any times that we observe these prices and the algorithm at that price, the algorithm updates the Q-value according to the formula that you find below the table, okay? So the cell that has not -- for the cells that have not been visited, there is no change. The algorithm updates one cell at a time. That is so by design, which implies that these algorithms are slow to learn, okay, because they update one cell at a time.

For the cell which is visited, the updating is, you know, with -- there is a way to 1A, which is given to the past value, so an alpha is a relative weight of new information, and the new information is the current -- the profit that the algorithm observes. By the way, when we are in this cell, the algorithm observes also what the other has charged, okay, and, therefore, observes the profit and the price of the opponent. Given the profit, it can calculate the value of the next state which is achieved and updates the Q-value according to the formula, okay?

Now, of course, for that to work, we have to instruct the algorithm to experiment. To allow the algorithm to learn, it has to experiment, which means it has to try actions that would not be optimal in the light of the information that he has acquired so far, okay, to learn something new. So there has to be an experimentation strategy, and there are different types of experimentation strategy.

What we use is the Epsilon grades. They're called Epsilon grade experimentation, which means that with probability 1E, in each period, the algorithm charges the price which, given the past prices, has the highest Q-value, okay? So in this row, you would look for the action with the highest Q-value, and the algorithm would charge that one, but with probability Epsilon, the probability randomizes uniformly across all other actions, okay?

And we have this probability decrease over time, starts from one. That is, initially, the algorithms randomize uniformly, and then, as time passes, they give more and more weight to the greedy -- it's called greedy action -- and less weight to the need for experimenting. So that is how the algorithms work.

The economic model is a standard model with

logit function, a fixed number of firms, constant margin of costs, okay? There are a bunch of parameters, and we do a robustness analysis with respect to all of these parameters. Okay, so that is what we -- the value of the parameters in our baseline experience.

M is the number of prices, okay? As I said, the action space has to be discretized, okay? So we start from 15, it's reasonably large, meaning that it is not easy to coordinate. If we had two prices only, for example, the game would become sort of a business dilemma. There is only one way to cooperate in a business dilemma. So you might argue that the problem is too easy for the algorithms. No, but with 15 prices, the problem becomes already quite difficult, and we also look at a much higher number of prices, up to 100.

Then these prices are supposed to range between -- somewhat below the Nash equilibrium price and somewhat above the monopoly price, okay? So that c-whatever is the gap between the lowest price we consider and the Nash price and the gap between the highest price that we consider and the monopoly price.

K is the memory, okay, so one-period memory in the baseline, and it's the number of players. By the

way, if we change these parameters, so if we enlarge the number of prices or we allow for a longer memory or for more players, one effect of that would be that the matrix would become bigger, okay? So in our baseline specification, the matrix is about 3300 entries, okay? And because each entry has to be visited a number of times in order to allow the algorithm to learn, okay, inevitably, it takes a long number of periods in order for the learning to be completed, and if you increase the complexity of the program, which translates into a bigger matrix, then the algorithms are inevitably even slower to learn, okay?

Then there is the discount factor, okay? It's pretty low, because you must consider that the interaction can be with algorithm quite frequent, okay? So an interest rate of 5 percent is not -- even on a yearly basis, and the period of time that we have in mind is much shorter than the year, but we choose this low discount factor to stress that that is not the limiting factor, as somebody claimed in the early -- in the early debate, okay?

All right. So what do we find? Well, first of all, there would be, in theory, what we know theoretical about this learning algorithm is that they are not -- they do not necessarily converge to

something, but we find they do converge, and here is the prices which they converge to.

There is a lot of noise, okay, because as I said, they learn by trial and error, and, therefore, the way they learn to cooperate is very specific to the pair of algorithms which are interacting, but overall we see the colored squares are the monopoly prices or the Nash prices. Prices are biased towards the monopoly price, and correspondingly, we have a profit gain. The profit gain is the ratio between the gain with respect to the Nash equilibrium profit and the maximum gain, okay? So it depends on the learning and experimentation parameter, but in the whole range that we consider, it's about 70 percent, okay?

To give you a sense [off microphone] experimentation parameter, the highest value that we consider, which corresponds to visiting a cell by chance only -- on average only four times over an infinite time horizon. So that would be actually too little to have decent learning, and actually, we then look at the representative experiment in which a site is visited by chance some 20 times, a point more or less here in this table.

Okay, so we have -- however, even if we vary these parameters, still there is quite a lot of

collusion. Now, that collusion is rational; that is, our algorithms play a Nash equilibrium. Of course, it is an infinitely (indiscernible) game, so there are many Nash equilibriums -- Nash equilibria, but what they do is to play one of these, okay? In our representative experiment, in 50 percent of the times, for each combination of parameter value, we run a thousand sessions, okay? In 50 percent of these, they do play a Nash equilibrium, okay?

In the remaining 50 percent, they come quite close to it; that is, even if they do not play Nash, what they could gain by playing Nash is less than 2 -- 0.2 percent more than what they gain in the actual play.

Another test of equilibrium play which actually makes a difference with the previous one or two is that when we design the experiment in such a way that the Nash equilibrium is -- the starting Nash equilibrium is the only equilibrium of the game, it is when data is zero, close to zero, or when there is no memory, our algorithms to learn to play the static Nash equilibrium, okay? So they do not charge high prices because they make mistake. They charge high prices because they learn rationally how to do so.

And actually, to verify this, we look at what

we call an impulse response. So once the algorithms have converged, have completed their learning, okay, which we take to mean that for 100,000 repetition in a row, they do not change their strategy, okay? Once they've done that, then we step in -- step in and force one algorithm to undercut the rival and see what happens, and that is what happens, okay? So this is the known deviating algorithms, okay?

So period one is where the deviation take place. The known deviating algorithm is caught off guard, so to speak, so it keeps playing the agreed-upon price, but then it punishes, and punishment is finite in length or width, and gradually they go back to the original point or cycle, and that is what the deviating algorithm do, okay?

So basically they do not play game trigger strategies. Why? Because there is a lot of noise due to experimentation. If they played game trigger strategies, okay, as soon as one algorithm by chance deviates, they would be stuck into a bad equilibrium, okay?

Actually, they learn to restart cooperation from any point, okay, even if we allow for bilateral deviation -- this is something that's not in the current version of the paper, but we are now working at

revision -- from any point, any possible state. Almost always they start to cooperate again.

Okay, I think that I've run out of time, so let me just tell you that the deviation -- of course, because the deviation is unprofitable, we ran bands of robustness checks, and the cooperation/collusion seems to be fairly robust, okay?

Now, just 30 seconds to tell what, in my opinion, is the main limit of this analysis. The problem is it takes a long time for the algorithms to learn to collude, okay? And it -- maybe in answering the question, I may provide more details about that, so what we should address next, and that is the topic that is most prominent in our research agenda, is what would algorithms that learn more quickly do, because there are algorithms -- true learning are simple.

We can't understand exactly how they work, but there are more sophisticated algorithms. There exist more sophisticated algorithms that are capable of learning more quickly, okay? So looking at these algorithms would be the next task in our agenda.

Thank you.

(Applause.)

MR. THOMAS: Okay. Thank you, Vincenzo.

Now, to discuss Vincenzo's paper, we have Wally

Mullin from George Washington University.

MR. MULLIN: Thanks. This is -- so this is a very interesting paper. So I think it's a real contribution, particularly on this link between the algorithms and tacit collusion. It's meant to be -- at least the version I looked at -- more about positive stuff than normative stuff, but obviously the normative stuff becomes important later, and for this environment, that's important, too.

So he uses experiments or numerical simulations to assess the algorithmic pricing. There are some assumptions he went over. So, for example, the basic demand structure is logit demand differentiated Bertrand safe game and repeated infinitely.

Second, also reinforcement learning, in particular Q-learning that he went over, and I think even that the verbal thing here, and/or his revision, the example -- because I'm not -- even though I do collusion, I don't know that much about algorithms, frankly, although I do worry about it from a social perspective, but the -- you know, going to the table, you know, the example I think was actually good, to keep, you know, the idea of but not, just (indiscernible) examples and framework examples in terms of the Q-learning.

We don't whether actual agents are using that, but on the other hand, it's computer science, you know, it's relatively simple, and there's maybe, like, you know, a reasonable number of parameters that have to do with -- you know, that have to consider readily economically interpretable -- and as I said, going through the exhibit, the Q matrix helped, okay?

The second or one of the other parameters is how much exploration takes place, and that's captured with this epsilon greedy model of exploration. So the greedy action is doing what the Q-algorithm would tell you to do, and then the Epsilon part, which could be small, like 10 on a, you know, scale of 0 to 1, would be doing something different, as he described it, kind of uniformly across the different examples, different choices, different aids, given whatever the state is. So that's an important feature, too, because -- and the final thing would be bound memory (indiscernible), and he did like one or two in the talk.

So what are the results? First, supercompetitive prices without communication, so that's obviously of potentially significant interest, because at least in the U.S. context that I'm most familiar with, like, there's a sort of vacuum, and obviously there's a much longer debate going back to

Posner and earlier, well, gee, should we treat tacit collusion as a normative matter differently than formal collusion, and obviously a part of this argument is, well, maybe it's about the evidence-securing process of formal collusion as opposed to what happens with tacit conclusion and/or the, you know, type one versus type two issue. I certainly don't want to punish someone who's being competitive because they're matching their rival, obviously.

And, you know, part of the argument as calculated in his book was going through in great detail about, well, what happens if you have two firms and they're doing -- using flags to communicate, but they're basically still agreeing, right? We would probably think that that would still be something legal and have it legal, okay?

So another part of this all is finite duration with gradual returns get cooperation, so that's also important in terms of, you know, the exploration model, because even as Epsilon gets smaller over time, which is going to occur in the limit, you'll still have this -- eventually you'll have -- ending up, as he says, you can have undercutting just because of the experimentation, so -- or the exploration. He wouldn't want to have grim strategy, which I understand.

So the first comment -- so certainly I know for this version you want to basically stick with a standard demand specification and then do what you want to do in terms of adding about the algorithmic pricing has on logit demand. However, I think because it is numerical, one possible thing for some other version would be to think about consumer demand. Countermeasures may be something strategic, right? Like, even if we're not in this formal collusion world, the consumer doesn't want to pay higher prices than they have to.

So what would happen if we had a procurement auction, obviously not with necessarily the Bertrand search, if consumers had commitment or they could move first or these sorts of things. So that may be more of a theoretical issue, but I would think that it would be not necessarily that hard to write the program that would do that.

So he actually -- in the talk, he talks about the policy stuff, which I think is actually very important. So one possible -- just kind of -- the U.S. context finds in damages, and then a second thing would be leniency or bounties. Now, leniency in terms an evidence course, because it's tacit collusion, there's not going to be this steering process, so how would you

really have a legacy -- a leniency, sorry, or a bounty in that context? Probably not.

However, competition policy players are going to care about it. So basically the other thing I would sort of take from leniency is only entities trying to make the incentive compatibility constraint for collusion harder, right? We don't want to -- we have to think through carefully, right, because we don't want to underrun the funds that we do have by giving leniency too readily.

But there's still this issue of, well, gee, if in a world of -- I'm not necessarily saying the competition policy authority is going to be running an algorithm, but if they can -- if they're aware of algorithms, one other thing would be get rid of them completely, right? But the other thing might be, well, you want to get -- now, for this paper because this paper is tacit solution, but it also sort of says here's the set of policy tools, at least here's a simulation of the -- numerical simulation of the policy tools a competition policy authority would want to implement in order to help make prices more competitive, take advantage of this incentive to undercut. That's about it.

MR. THOMAS: All right. Thank you, Wally.

So now I will invite Vincenzo back up to the podium. We have about 12 minutes for questions from the audience. So please just go ahead and raise your hand if you have a question.

Is there -- yeah, there you go.

AUDIENCE: Hi. Thank you so much for your presentation. It was a really interesting paper.

So my question is if we were to figure out a way to sort of regulate this and sort of maybe, like, prevent this, what is the strategy? Do we, ex ante, say to firms you can only do certain parameter values that -- you know, like you showed K equals a zero, shuts down the -- certain collusion, right?

And do we say them to them, ex ante, you know, you can't have these types of algorithms or these types of parameters, or do we, ex post, come in and say, look, we observed the prices that you ended up charging, and we're going to somehow, like, take money back and refund it to the consumers, and then hopefully the firms will build that into the algorithm and then the algorithms will learn to take this regulation into account? Thank you.

MR. DENICOLO: Well, thank you for the comments. Actually, this issue of the randomness was also raised by the discussant, by Wally. Personally, I

would -- well, as I said at the beginning, whether you would attempt to regulate this industry or resort to export intervention, antitrust, may depend on your assessment of the likelihood that the algorithms may collude. If you really believe that collusion is inevitable, then you might go for a sort of ex ante regulation as has been proposed by some scholars.

And in that case -- but even if you intervene ex post -- and I'm more in favor of this second policy avenue -- there is still a problem of remedies, okay? So basically one thing that I did not say at the beginning -- but it's definitely true -- is that, you know, these pricing algorithms are good from many points of view. They allow firms to react quickly to market conditions, and in many cases, that may have pro-efficiency effects, okay? So we must be careful not to throw the baby out with the bath water when we regulate or we prohibit these practices.

I think that part of the research which is necessary in this field should be precisely devoted, once we have established that there is a risk, and some policies should be doing something, which I believe is the case, how can we modify the algorithm or restrict the algorithms in such a way that they will not be able to collude?

So our results show quite clearly that if, for example, we do not allow an algorithm to condition its current price on the past prices, then that would stop collusion, okay? But, of course, that remedy is quite strong because the past price of the rival could be regarded as part of the environment.

Actually, in this framework, it is not. In this framework, the past price is payoff-irrelevant, okay? If you look for a mark of strategies, that should not depend on the past price, but there could be many reasons, okay, many dynamic effects, which implies that the past price also should be something that you can condition on.

So it's not an easy problem, but for the moment we are asking kind of preliminary question. The question is, should we worry at all or should we not, because there are people who argue that we shouldn't worry.

AUDIENCE: I'm over here to your right.

So my question is, the firm's choice of which algorithm to use is in the real world an endogenous choice, and so you might think that they could choose among a whole variety of algorithms, and I'd be interested to know -- you know, that makes the strategy space huge if you think about all the algorithms they

could choose among, but I wondered if you had thought at all about "endogenizing" the choice.

Like, for instance, you could stick within your model and just allow the firms to choose the data independently in order to maximize their utility, or they could pick the Epsilon independently and, you know, tweak the algorithm independently from one another in a way that maximizes profits. And I'd be interested to know, do you still get elusive outcomes when they're -- when that choice is endogenous?

MR. DENICOLO: Well, thank you for the question. That's very interesting. Actually, we just, you know, started to look at these kinds of issues. So, for example, in one of our robustness analyses, we allow for -- we only look at Q-learning algorithms, but as I said, our plan is to look at quicker ones in the future.

But within Q-learning algorithms, you can change the parameters, okay? There are two key parameters, the learning rate and the experimentation rate, and in addition the way the matrix is initialized. Now, we tried -- we look at what happens when the algorithms have different learning or experimentation parameters and who gains most in case these parameters differ, which is kind of preliminary

analysis towards asking the sort of questions that you raise. It's definitely very interesting but something that we haven't done yet.

AUDIENCE: Hello. Hi. Am I on?

This actually is a followup question. I didn't know it. Imagine your situation in which you allow in learning algorithms a choice of the rate. My intuition is that competitors would want to make their algorithm more complex. That's one question. So if you test out your experiment, will you find that there are gains to be more complex than your rival?

The second is that you sold your results as saying even with simple algorithms, you get collusive prices. It isn't obvious to me that with complex algorithms you would still get collusive outcomes. So the rat race to have more and more complicated algorithms could lead to simply dead weight loss associated with rating faster, if you will, or it could conceivably change the competitive aspects of the market. I wonder if you have insight in that.

MR. DENICOLO: Well, thank you very much for the question.

Actually, one thing that we say in the current version of the model -- of the paper, sorry -- is precisely that more complex algorithms would be more

likely to collude, but we are deleting this kind of speculative comments from the paper on the request of the editor, and I think -- I mean, if I were him, I would have asked for the same, and we just limit ourself to the report that is out. So, I mean, only by doing more simulation and more experimentation can we really answer to the question.

Let me, however, concerning the issue of more complex algorithms being able to exploit the current one, the current simple ones, well, actually, at least upon conversion -- convergence, so once the learning is completed, as I said, our algorithms play a Nash equilibrium or something which is pretty close to a Nash equilibrium, so they cannot be exploited.

Now, what happens during the learning? That is really very, very difficult to assess unless basically you do the analysis. So I will refrain from engaging in further speculation and simply answer to your question saying that, well, this is something that will have to be seen by conducting further analysis.

AUDIENCE: Hi. Thanks again for the paper and the presentation. I'm right here.

You had done an analysis in which you looked at the response to deviations, and it looks like whether they are unilateral or bilateral deviations, you find

that the algorithms converge back to the collusive outcome. So in that context, I was wondering if you had looked at efficiencies in pass-through and what the implications are of that result for what's likely to happen if one firm experiences efficiencies. Is it likely to just retain the efficiencies or pass any of it on?

MR. DENICOLO: No. To be honest, we haven't looked at that. So what you have in mind is a kind of situation where, in the course of the interplay between the firms, some of the parameters which characterize a firm changes, and you -- well, that -- that's a good suggestion for future work, but we haven't looked at that.

AUDIENCE: One -- one -- I'm over here.

MR. THOMAS: I'm sorry. I think we actually need to move on, but thank you. Maybe you can chat after the -- after -- during the break. Thank you again. Thank you, Vincenzo.

MR. DENICOLO: Thank you, and thank you for the questions.

(Applause.)

MR. THOMAS: So next we have Alminas Zaldokas from the Hong Kong University of Science and Technology presenting a paper titled "Corporate disclosure as a

tacit coordination mechanism: evidence from cartel enforcement regulations." Thank you.

MR. ZALDOKAS: Good morning, everyone, and thanks for including this paper on the program. How do I move -- how do I move -- the green one? Okay.

So thanks for including this paper in the program. This is joint work with Thomas Bourveau, who is at Columbia Business School, and Guoman She, who is our Ph.D. student at HKUST. (Indiscernible) presented it yesterday, so apologies for that.

So in this paper we're looking to the -- empirically looking to the corporate disclosure as one of the mechanisms, how tacit collusion can be sustained in the product markets, and we kind of started looking into this case after we stumbled into this one particular FTC case. So let me remind you briefly of that.

So U-Haul was holding an earnings conference call. So what is earnings conference call? Basically after the firms announce annual earnings or quarterly earnings, the representatives of the firm -- CEO, CFO, other top management -- often hold a conference call with equity analysts.

So they're holding this conference call on February 7, 2008, and they knew that the

representatives from the main competitor, Budget, are also dialed into that call. You know, these calls are often accessible to all potential investors and the public.

So CEO of U-Haul makes the following points. U-Haul is acting as the industry price leader. The company has recently raised its rates, and competitors should do the same. To date, Budget has not matched U-Haul's higher rates. This is unfortunate for the entire industry, and U-Haul will wait a while longer for Budget to respond appropriately. Otherwise, it will drop its rates.

Okay, so this is like verbatim quotes from the conference call transcript, describing pretty much, like, you know, (indiscernible) strategy, if you want. So FTC, to our knowledge, has started investigation, and let me just briefly summarize that it basically was trying to say that this information is not necessarily targeted at investors but might be targeted to the main rival, Budget.

And further investigation actually discovered an internal memo from the CEO that Budget continues in some markets to undercut us on one-way rates. Either get below them or go up to a fair rate. Whatever they do, let Budget know.

So this is -- this is probably -- the second part, which probably can be interpreted even as explicit collusion, formal collusion, and the first probably is more tacit because it's using a public communication rather than a private meeting that the second -- the second strategy would do.

So what we -- there was a similar case with Valassis Communications and used pretty much very similar wording, so let me not go through that, but what we tried to think about in this paper is whether this can be seen as a more regular pattern in the data, and in general whether, you know, communication that is primarily targeted to investors, that we will call -- will be calling financial disclosure just for the sake of being short, is actually -- can be used as a sort of tacit collusion stabilizing mechanism.

Following these two cases, you actually can see that some of the lawyers have started looking into this, and there's a presentation from one of the law firms that says, you know, what are the potential strategies to avoid antitrust traps, investor analyst calls, and I just have one slide from their presentation.

It says, you know, why are private plaintiffs listening? Why is the government listening? And they

discuss the theories of antitrust liability and so on. And they specifically mention that this -- that certain -- certain discussions in the conference calls can be seen as the invitations to collude and unlawful signaling, as this U-Haul case has actually shown.

So what we will do in this paper, we will ask do firms use financial disclosure to share information that could benefit peers in these tacit collusion arrangements. We will not look into all disclosure that a firm is using. We will look into two particular types of disclosure, and the first one is this particular cases that I mentioned, is the conference calls with equity analysts, and the second is the material contracts with customers.

I will go into describing these particular types of disclosure a little bit later. Before that, I will describe our identification strategy, but in short, the first one will be the public communication with investors, and the second will be the contracts that the firm has to submit to the SEC. If this is a material contract that poses a significant risk to the investors, they have to submit it to the SEC as SEC filings, and they can choose whether to redact certain information from the contracts or not, and, you know, in particular, what we would be looking into is whether

they redact product prices or they disclose them publicly.

Now, what is the tradeoff that we have in mind, that the firms are facing? So we're all in finance and accounting, so from our perspective, as the finance academics, you know, we have been learning and we have been kind of as a first order effect agreeing that more transparency brings benefits of reduced information asymmetry, so firms are actually inclined to produce more information to the public.

This leads to less adverse selection, leads to lower cost of capital, better governance in the case of the moral hazard. So, you know, as a first order effect, finance accounting really likes more information provision. This helps both firms and investors make more optimal decisions.

However, if you start thinking that this information can also be used strategically by other firms to tacitly coordinate the actions in public markets, maybe we should actually be considering also that there could be some welfare costs of disclosure.

So our goal with this paper is pretty much to give an example to particularly finance researchers of how these first order effects of higher disclosure leading to lower cost of capital might have negative

welfare consequences because of this potential conflict between securities and antitrust regulations. Yeah, so securities regulations would ask for more disclosure, and antitrust regulations might actually try to limit disclosure if this disclosure leads to negative welfare consequences from the, you know, consumer point of view.

Now, this actually is being discussed in the legal literature as well, and in some sense it goes back to the Supreme Court ruling in the Credit Suisse vs. Billing case. That case is not related to what we're doing in this paper. It was about an IPO -- investment banks and IPO pricing, but one of the outcomes of that legal case was the Supreme Court's kind of suggestion that creation of the SEC implicitly exempted the regulatory securities issues from antitrust loss, and then the question that arises in these recent legal papers is whether, you know, the corporate disclosures that are targeted to investors are, indeed, immune entirely from antitrust challenge under this Supreme Court decision and, for instance, whether investor calls or whether the SEC filings are considered as the -- you know, as part of the regulatory securities industry and so antitrust laws should not be applying there.

So what we'll do in this paper, we will apply a sort of reduced-form empirical strategy to study whether there's actually the effect -- sorry, whether there's actually a prevalence of -- that we can say that under certain conditions firms turn to using financial disclosure for tacit collusion. So, you know, as an empiricist, I'm -- you know, we are trying to understand whether -- we are trying to -- ideally trying to find the duration in tacit collusion.

Tacit collusion obviously is impossible to observe for the empirical researchers based on the observable data, and so we'll try to come up with an empirical strategy that would give us evaluation in the incentives to tacitly collude.

So we will make this identifying assumption that when explicit collusion costs increase, when it's harder for the firms to meet privately and discuss prices, at least for some firms, tacit collusion becomes the most appealing alternative strategy. So, of course, you know, if explicit collusion costs increase when the antitrust enforcement is higher, some firms will turn into, you know, competitive Nash equilibrium. Other firms might -- we argue might actually find other ways that are potentially collusive through informal means to maintain some sort of market

power.

So what will be our identification? We will look into the strengthening antitrust enforcement around the world that affects U.S. firms and presumably increases explicit collusion costs. So whatever I will show in this paper will be a joint test of antitrust enforcement against explicit collusion, making tacit collusion more likely, and tacit collusion changing financial disclosure in product markets. So, of course, I will not be able to separate the two effects.

What we don't do in this paper is we don't say that this is the only way of how the firms will be adjusting and that we only wanted to say that this is a kind of -- a possible way of how the firms could be adjusting to the increased antitrust enforcement.

Of course, they could be doing something else. They could be merging. They could be colluding by some other disclosure, for instance, advertising. They could start sharing what's now popular to study and discuss in the, you know, policy debates, common owners, common lenders, you know, common other partners, common board members.

So what we are only trying to do with this paper as a sort of disclosure on our behalf is that empirically document that there's a potential conflict

between antitrust and securities regulation and, you know, provide this -- these findings that we have for further debate on the normative implications, so...

All right. So let me talk about the identification strategy. We'll look into the leniency programs -- I'll call them leniency laws -- that, you know, were adopted in a staggered fashion around the world. As a short reminder, these leniency programs grant immunity to the first self-reporting cartel member and allow for reduced sentences to other cooperating members. In the U.S., that also includes criminal liability for individuals in the firm. Most other countries, there is limited criminal liability, but -- and this is not something that we study anyways.

So I have a trilogy of papers that look into this global enforcement of antitrust through leniency programs, and in the first paper that just came out in RAND on this issue, we find that the number of convicted cartels increased by 154 percent after this leniency program's adoption around the world, and the gross margins of the publicly listed firms dropped by more -- you know, more than 14 percent.

We -- in this paper, in these earlier papers, we collect data on the timing of adoption of these laws, so in the -- for the U.S., we used the

strengthening of the law rather than the adoption in 1993, and for the other countries, we look into the most relevant data that we think the law has started to be the most -- the most welcoming to the -- to the cartel members to be used in talking to antitrust agencies and the judicial agencies.

So we collected data on 63 countries and territories. For some countries and territories, we could -- there was no law at the time of our study period from 1990 to 2012, and they are listed in this yellow square. So Hong Kong, for instance, at that time didn't have any competition policy, and now they have just adopted it pretty recently.

Now, one particular concern of using this staggered adoption of laws, any laws, is that there is some sort of underlying economic reason why they're being adopted, and this underlying reason might actually be driving whatever outcome we're looking at. So what we try to do is to read through the media documents in local languages to see what is the media discussion, what is the policy discussion around the adoption of these laws, and we don't have -- of course, for every law adoption, there's a particular reason -- a particular political economy reason why it has been adopted, but we don't find that there's a single

underlying correlated reason.

So, for instance, U.S., Switzerland, Hungary, laws were passed after significant collusion cases. So Hungary, I think, telecommunications; Switzerland and U.S. was pharmaceuticals. Taiwan was a concern about rising consumer prices, so, you know, maybe some good economic conditions. Korea was opposite. There was a financial crisis, bad economic conditions. Mexico was, you know, suggested by OECD to adopt it.

Singapore was pushed by U.S. using -- during the negotiations on the free trade agreement. Some EU member states were -- got pressure from EU. IMF/World Bank sometimes asks for the overhaul of antitrust laws as part of their funding. So there's some different reasons, but there's no one reason that we actually can see is happening.

Now, in this paper, we're looking into the U.S. firms. So the way we'll do -- we'll apply this identification strategy, we look into the staggered passage of laws in the countries with which the firm's industry trades. Why would it matter, the passage of laws in other countries? Well, it -- in our -- to our understanding, the coordination between antitrust authorities becomes easier. Even if they don't coordinate, even if they are -- a leniency applicant

asks for the information not to be shared with the other antitrust authorities, which is often the case, the information on foreign cartel often becomes public, which helps U.S. authorities to prosecute similar behavior by the same firms domestically.

Even if they don't do that, there could be private civil cases in the U.S., and here, of course, we can remember the chocolate cartel that led to the private cases in the U.S. based on the antitrust enforcement in Canada and European Union.

In any case, it makes it more difficult to form international cartels with industry peers, and we think this is kind of somewhat exogenous to the economic conditions surrounding the firm back in the U.S. So it is not driven by the industry effects that could be a concern here in this identification strategy.

To be, you know, more specific, here are a few quotes from DOJ. "The viability of foreign leniency programs is critical to U.S. antitrust enforcement efforts. Over 90 percent of international cartels that have been prosecuted by the Division were active in Europe as well as in the U.S. The U.S. leniency program served as a model for leniency programs adopted in other jurisdictions, and U.S. is now almost always joint investigating and punishing international cartels

by European Commission, Japan, Brazil, Canada, Australia, and others."

So I think this cooperation with the other antitrust agencies kind of suggests that this passage of the international leniency programs is important also for the U.S. firms. In particular, it will create this treatment measure, which we will proxy from its exposure to the passage of foreign leniency firms, and will use the firm's industry imports as the sort of weighting variable.

So we will create this weighted average of foreign leniency laws where the weight will be the share of the firm's industry's imports from certain country to U.S. as a fraction of the total industry's output.

So the separation will be the country in the severe level, so here's an example. When Spain passes the law in 2008, foreign leniency increased by X percent for U.S. industries that import X percent of total output from Spain. That's how, you know, this measure develops across different two-digit SIC industries.

And then we'll do sort of essentially a difference-in- difference estimate, where we use the two disclosure variables that I just -- that I will

introduce in a second, and control for the firm and year of fixed effects, and use this foreign leniency measure that will make some firms -- some industries more treated in a year, that trade more with countries that pass leniency law, and other industries will be controlled that trade less with countries that pass leniency law in that year.

All right. Just as a validation, we see that this foreign leniency variable actually leads to more cartel convictions in the U.S., leads to lower gross margins for U.S. firms, lower stock returns, and lower Producer Price Index at the industry level. So, you know, there's some effect validating that it is an important measure or important -- it's important to look into the foreign antitrust enforcement for U.S. firms. They react.

Okay. So which financial disclosure measures will we be looking at? Again, this is -- these are material contracts with customers and discussion of competitive strategies in the conference calls with equity analysts. Let me describe the first one. So disclosure -- the first disclosure method that we'll look at is the new material contract with customers. Regulation S-K requires firms to file all material contracts, let's say contracts with the CEOs, contracts

with the board members, contracts with suppliers, something that has some material risk for the investors. So we think this is a credible forward-looking measure about prices and quantities if we look into the contracts with the customers, okay?

And this, of course, it can't be cheap talk, because these are actual contracts, and the firms will be liable if they provide false information. So due to propriety reasons, the SEC allows firms to request part of the information to be withheld from the filings, and firms have discretion whether to redact some information or not. So we kind of measure whether the firms actually redact information by searching whether they ask for these confidential requests or not. And actually, when the firms ask for the confidential requests, SEC almost always grants them. So 98 percent of the cases, they grant them.

Here are two examples. There is one redacted contract. MIP agrees to supply products to Biomedica, and we see there's prices that are not disclosed for the investors, okay? And here's the example of the unredacted contract. There's another industry, but we see that there are prices that are disclosed to everyone, including the firm's rivals, to see.

Now, the second disclosure method that we'll

use is the conference calls that I mentioned. We concentrate on the CEO's and CFO's presentations, and we just -- first, we count product market related words. Then we look into the 20 most frequent words in the two earnings conference calls that led to FTC cases. Then we just look into the -- whether competitors are mentioned during conference calls. We also do a little bit of machine learning to try to enhance this analysis. So we have a number of dictionaries that we create to measure, you know, how much the firms disclose of product market strategies during their conference calls.

All right. So this is pretty much the main table that I wanted to show. When this foreign antitrust enforcement becomes stronger, we say presumably from some incentives to switch to the tacit coordination, and they do that through the financial disclosure. We see that the contract redaction decreases, so that is they disclose more prices, more quantities in their product market contracts that they submit to the SEC, and they disclose more of the product market strategies during their conference calls. So this is just sort of the main take-away from the paper.

We look into whether there is a pretense in

that. We do not -- we do not really find that. Most of the effect comes in the first few years. And this is an important slide to -- so that I want to spend a little bit more time. Most of the effects -- most of the sort of validity of this argument and the identification strategy I think should come from the -- seeing whether the effects are stronger sort of in terms of the cross-sectional implications where we expect them to be strong according to the theory, and we'll look into three steps here when we explore the heterogeneity of the effect.

First we will look into the industry's propensity to engage in collusive arrangements, so we'll look into whether it -- the effects are stronger in more concentrated industries. In the industries that we can say has more homogenous products, in the industries that have more recent or more -- more recent convictions or more convictions that can be predicted from the observable data, more usage of patents or less usage of patents, actually, in this case, and higher industry growth.

So, of course, all these predictions have been also argued in theoretical literature. Probably some of these at least could be seen as opposite, but I think it's kind of comforting that at least as a first

order effect we see the effects stronger where it is.

Second, we see that the industry's ability -- we look at the industry's ability to sustain tacit coordination by unilateral disclosure, how many public firms are in the industry, and whether these are, like, strategic complement versus strategic substitutes industries. So we, indeed, see that the effect is stronger for the industries that have more public firms, so there is more information that is being provided, and also for the more specific complement industries. And, finally, for the firms that are larger, we see a larger effect.

We do like a bunch of robustness checks. I will not go through all of them. I just want to mention that we're look into firm-level identification based on where their subsidiaries are. We look into the countries with higher judicial enforcement versus lower judicial enforcement, and -- and we -- one thing that I want to mention, we don't find effect if we look into supplier contracts rather than the customer contracts.

There are a few sort of alternative responses we can study, maybe the effect that we have is actually just the firms are responding to more competitive environment. Maybe they want to raise equity to

compete with peers, so they increase disclosure. Maybe they want to signal good behavior to antitrust authorities and, thus, try to reduce litigation risk, or maybe the effect is confounded with alternative coordination channels, such as the increased public advertising, they coordinate through common ownership.

So what we try to do is actually to exclude certain subsamples where we see most of the increase in the competition, most of the increase in the advertising, in the common ownership, and we see that our effect prevails. And as a last thing that I want to show is to see whether this change in disclosure is consistent with tacit conclusion; that is, if the firms adjust the disclosure, whether they see better outcomes in the product markets, and we see that the firms that do not increase disclosure after the passage of the leniency laws, they see a decrease in gross profit margins, and the firms that actually increase disclosure during the product -- after the leniency laws actually see no change in the gross profit margins or just a smaller decrease.

So let me conclude here. Basically we tried to argue that the increase in explicit collusion costs leads to a strategic change in firms' disclosure, and the firms try to produce more information in their

financial documents on product markets-related matters, and this is consistent with financial disclosure being used as a coordination mechanism to sustain tacit collusion.

We think this has some policy implications about a conflict between the securities and antitrust regulations, and there's some OECD discussions on this that I think also raises a similar concern.

I look forward to the discussant, and thanks so much in advance for your comments.

(Applause.)

MR. THOMAS: Thank you, Alminas.

Now, to discuss Alminas' papers, we have Leslie Marx from Duke University.

MS. MARX: Thanks to the Federal Trade Commission for hosting the event and for the organizers for putting everything together, and I appreciate the opportunity to contribute. I'm getting the "hold on a second" signal from the tech guys.

I've been trying to get my accounting colleagues at Duke interested in issues related to collusion for years, and so I'm very excited to see accounting professors taking a look at the issues related to collusion. I think the SEC regulation versus antitrust investigation contrast is really

important and the different views on transparency. Super.

This paper analyzes two key types of data of interest for studying collusion, earnings call transcripts and SEC filings. In particular, the authors analyzed firms' strategic use of financial disclosures to sustain collusion. The conduct that they focus on is increased product-related discussions in earnings calls and decreased requests for confidential treatment on material sales contracts in their SEC filings.

Their high-level interpretation of the results is that the passage of leniency laws in countries outside the United States has caused U.S. firms to switch from explicit to tacit collusion, with the tacit collusion being based on communication through the earnings calls and SEC filings.

Since the authors focus on earnings calls and SEC filings, let me start by thinking about how those might facilitate collusion. First, they might help to coordinate future conduct in the same way that public price announcements might be used by colluding firms. Second, they might help to monitor past conduct if SEC oversight can be relied upon to ensure truthful reporting. And third, they might be used to mitigate

buyer resistance. For example, firms might negotiate a particularly aggressive contract and then deliberately reveal those contract terms through their filing in order to provide a beneficial anchor for future negotiations.

I am completely willing to believe that firms will use all legal and likely some illegal avenues available to them to enhance profits, including communication and signaling via earnings calls and including communication and signaling via SEC filings. So I think this paper makes an important contribution in taking a look at these things with an eye towards their anticompetitive implications. That said, the research also raises questions about alternative explanations for the results, and I'll talk some about that later.

Let's look at the details of the results. The data is SEC redaction data from 2000 to 2012 and earnings call data from, I think, 1994 to 2012. The authors attempt to explain earnings call comments and the failure to redact contracts in SEC filings based on the introduction of antitrust leniency in countries that have exports into a firm's two-digit SIC code. They find that there's a causal impact of increases in explicit collusion costs where those are going to be

defined in terms of the foreign leniency policies on firms' disclosure choices.

Treatment variable is a measure of foreign leniency, so if J is the two-digit SIC industry, T is the year, and K is the foreign country, then the treatment variable for industry J and year T is the weighted sum of indicator variables for foreign countries having antitrust leniency in a given year, where the weights are the shares of the SIC code output inputted into the United States from that country in 1990. So, roughly, the treatment variable is the import share and the two-digit SIC code from countries with antitrust leniency.

The authors provide a graph -- this was shown -- of the treatment variable for the various two-digit SIC codes, and I wanted just to highlight a couple of these. So in this PowerPointed version of their graph, I highlight two of them. So the red line in this graph is SIC code 28, which is chemical and allied products, and so it's the red line there. You can see it has a jump up in 2000. So these are chemicals, and Germany is a big manufacturer. So 2000 is when Germany instituted antitrust leniency, and others it's basically a constant.

And then the blue line there is SIC code 36,

which is electronic and other electric equipment. That has the same jump in 2000. Also, the next jump comes in 2005. That's when Japan instituted antitrust leniency, the electronics manufacture. And then in 2006, again there's a jump, and I'm guessing that relates to Mexico. They have significant imports to the United States as well, and they introduced leniency in 2006. Overall, the blue line is more or less an upward trend there.

And it is noted at the top on this slide, if you use leniency at the European Commission level, which they do in robustness checks, that moves that 2000 bump up from Germany over two years, to 2002. So the red line would have the change in 2002 instead of 2000.

The disclosure variable for contracts is a 0/1 variable that's coded 1 if one set of key words -- one of a set of key words related to confidentiality appears in the SEC filing exhibit, so things like confidential treatment, confidential request, or confidential redacted. The earnings call variable is a share of -- so a number between 0 and 1, a share of words in the CEO or CFO's opening statement that are related to certain key words, like product, service, consumer, customer, user, or client.

Then there are other variables included, the return on assets, total assets, the HHI for the two-digit SIC code, and import penetration at the four-digit SIC level.

Turning to findings, the authors find that the passage of leniency laws leads to a dissolution of cartels, and here they're regressing the number of convicted cartels on the foreign leniency variable and finding significant effects. This would, indeed, be great news for the role of leniency programs, but I think we should be careful before reaching strong conclusions based on this particular result. In particular, we don't know whether leniency laws have supported the initiation of other cartels.

In additional results, the authors regress the redacted contracts, that 5/1 variable, on foreign leniency and find statistically significant reduction in redactions, and they regress the percent product words from conference calls on the foreign leniency variable and find a significantly significant -- statistically significant increase in the use of product words. So that's interesting.

So then the question is, what's going on? So these effects are showing up here, so let's try to think about what might be happening here. The authors

offer the following interpretation: Firms are explicitly colluding at time zero. Leniency abroad is bad news for sustaining collusion, so those firms switch from explicit collusion to tacit collusion, and they structure their tacit collusion around having fewer redactions and more use of product words.

So that may be happening, but I think more work is probably required to understand these results, because I think the results are also consistent with other interpretations, and I offer one here, and I'm not saying that I necessarily think it's what's going on, but I worry that it also can't be ruled out.

So in this alternative, firms are not explicitly colluding to begin with, and then the leniency abroad is good news for colluders, either because it lowers the expected fines if you get caught or the adoption of leniency itself is a reaction to some conditions that made collusion more likely or more profitable.

So then firms start explicitly colluding, and they introduce monitoring and enforcement structures based on the SEC filings and work to reduce (indiscernible) through both those SEC filings and the statements in the earning calls. So I think it would be interesting to dig into some more of the details

about what's going on behind those empirical results, and something that I think would be really super interesting in this setting would be -- and maybe not part of this paper, but maybe for future work -- would be to look at some case studies.

So what happens when you drill down to a particular industry or a particular firm and who are the firms that, in the author's interpretation, were explicitly colluding, because maybe we can prosecute them. So it would be interesting to look at some of the individual cases underlying this data and also perhaps take a more -- maybe a more nuanced approach to the unredactions as possible -- in some sense in which you can identify the intent behind the words being used in the earnings calls.

I think, in conclusion, it's really great to see these data being examined with an eye towards antitrust, and I applaud the authors' work collecting and analyzing these data, and I would be really happy if leniency programs in other countries were helping to solve antitrust problems here in the United States, but I still have concerns that something else might be going on. So I look forward to future work on this as we learn more about the details of the drivers for these really interesting results. Thank you.

(Applause.)

MR. THOMAS: All right. Thank you, Leslie. Now we have some time for questions for Alminas. Please just raise your hand and someone will bring you a microphone.

AUDIENCE: Very interesting results. I wonder whether your empirical facts could be explained by changes in SEC regulation, like, for example, if SEC is pushing for more disclosure or make it sort of less attractive to ask for redaction or confidential treatment?

MR. ZALDOKAS: So our identification comes from the exposure of different industries to foreign laws through their import shares or export shares -- doesn't really matter how we measure it -- so if this is driven by the SEC regulations or any other regulations in the U.S., our findings kind of should be correlated -- those other regulations should be correlated in terms of exposure with those import/export shares across industries.

So we thought about it. We looked into a few of the SEC -- changes in SEC disclosure, but it's hard for us to think of why those particular disclosures that affect all firms in the U.S. should have different exposures according to the trade for the different

countries around the world. So I think as long as we control for the industry trends in general or sort of year fixed effects and so on, the particular regulations in the U.S., it's very hard to believe that they explain our results.

AUDIENCE: You just said that it doesn't -- your results -- your empirical results are robust to how you weight, that is, the exposure, the bar tick kind of --

MR. ZALDOKAS: Yeah.

AUDIENCE: -- so the weight you give. Now, you used the share of -- you used the import share of -- into the U.S., right?

MR. ZALDOKAS: Right.

THE AUDIENCE: But you just suggested that if you use the export share, it doesn't matter. Can you just chat briefly about why you would use one versus the other?

MR. ZALDOKAS: Right.

AUDIENCE: It seems to me it would depend -- I mean, what would be a more natural approach to me would be to use the share of the firms -- or the industry, if you want -- the U.S. industry's value of shipments that they export, but it depends on whether you think this is facilitating tacit collusion among U.S. firms in the

U.S. or tacit collusion between the U.S. firms and foreign firms --

MR. ZALDOKAS: Correct.

AUDIENCE: -- abroad.

MR. ZALDOKAS: Right. So we thought about it a lot, and we also got comments from the editors on this. We decided to go with the imports for the following reason, that we think that since this is the communication that is -- we think is primarily targeting U.S. domestic markets, we're thinking about the potential collusion that is happening between U.S. firms and foreign firms in the U.S., rather than the potential collusion that is happening within U.S. firms and foreign firms outside of U.S.

So as long as these are -- these similar types of behavior are happening inside and outside, the foreign laws would matter to stop the collusion in the U.S. as well.

Okay. I would like to thank Leslie for the great comments as well, and I think -- I think the suggestion that the explicit collusion might be increasing and might be sort of used together with the tacit coordination -- or sort of -- it can be used as enforcement mechanism, I think it's a great thing to think about, and in general, I think that that is in

line with our arguments that we should be looking into this disclosure more carefully, even if -- and especially even more worrying if what you're saying is happening is in place, then if we should even put more efforts in trying to understand where this investor-based disclosure is actually only good for reducing information asymmetry and lowering cost of capital, or there may be some negative effects.

Thanks a lot.

(Applause.)

MR. THOMAS: All right. Thank you, Alminas.

Now we have a break, about 35 minutes. Please be back, seated, around 11:20. Thank you.

(A brief recess was taken.)

MR. THOMAS: All right. Welcome back, everybody, from the break.

So now it's my pleasure to introduce Joel Sobel, one of the scientific committee members for this conference. Joel is a Professor of Economics at UC San Diego who specializes in microeconomic theory, game theory, and information economics. Joel has held various editorial positions, including editor of Econometrica and co-editor of The American Economic Review, and Joel has received numerous grants and honors, including a Guggenheim fellowship and a Sloan

Foundation fellowship.

Joel has developed pioneering theories of communication, lying, and deception, which are directly relevant to the consumer protection portion of the FTC's mission. In addition, Joel has written on myriad other topics, including altruism, evolution, and optimal algorithms for counting the large numbers. His keynote address is titled, "Deception: Theoretical Considerations."

Please join me in welcoming Joel Sobel. Thanks.

(Applause.)

MR. SOBEL: There's something in red letters that says "Hello" here, so hello. There are many people that I have to thank, but the organizers from the FTC have been generous in their time, in their patience with me, and their support, both emotional and intellectual, so I'm honored to be here.

This is sort of a stretch for me, and the talk is going to be different in a couple of ways. One way is that I put a little bit on every slide, so I think in the bottom there might be some enormous number of slides, but don't worry. It's less than that.

The other is that because I was nervous, I inserted jokes, and prepared jokes aren't really good.

So I apologize when you're cringing. One of them comes early.

Which -- what goes forward, green?

MR. THOMAS: Green. It's the big green arrow.

MR. SOBEL: Green arrow. Thank you.

So this is an apology. I'm sort of -- I'm not sure why I'm here, but I'm a game theorist, and when you talk, this cartoon comes to mind. So I don't know if you can see the pictures. This is a famous Gary Larsen cartoon on the top panel. The guy is saying to the dog, "Oh, Ginger, I've had it. Stay out of the garbage. Do you understand, Ginger? Don't do this again." Then Ginger's listening to this, and it goes, "blah-blah-blah-blah," Ginger, "Blah-blah-blah-blah-blah-blah-blah." So, you know, I'm listening to you do serious economics, and I'm hearing da-da-da -- Nash equilibrium -- da-da-da.

(Laughter.)

MR. SOBEL: So that's enough about me, but I'm going to talk for a while, and my fear is that you may be bored, so it will be over in 27 minutes and 55 seconds.

Here's a roadmap. I'm going to tell you what got me thinking about the issues that I'm going to describe. I'm going to do this in a formal model. I'm

unapologetically a theorist. And I'm going to supply a definition of deception. So the motivation will tell me why I needed one. And I'm going to briefly say something about the difference between deception and lying. I won't define "lying." I can do that, but not today.

And then I'm going to try to indicate two things that are I think important ideas but intellectually kind of shallow ideas. So sometimes it's useful to say things that are obvious. I am going to contrast omission and deception because I think that's a sort of intuitive, interesting thing that you can say. I'm going to introduce an idea called "damage," and I am going to connect that to deception, so that doesn't mean anything now, and -- I'm not really going to do that.

So this is my motivation. I study strategic communication, and sometimes glibly I say, "I think about lying," and it dawned on me that I didn't really know what lying was, and then it dawned on me that I didn't know what a lot of other things were, including deception, and I was curious to see if I could tell a coherent story. It turns out -- and it's part of a broader research program -- I think that it's natural to observe that lying has costs that standard economic

models don't include and that this cost of lying depends on things.

In the background, some of the intellectual boundaries that I touch upon help identify when people have costs of lying. This matters. This is interesting. It's not today. So today maybe I'm going to do this, but it's going to be in the background. So be generous, and there will be a few comments. So this is my apology. It's be generous.

Okay, model. So this is a formal model. There will be seven or eight lines. It's fairly simple, I hope. So I am just going to talk about communication involving one player who knows something and another player who maybe takes an action. The sender knows data. The sender sends a message. Maybe the sender can take an action, too.

When you think deception, deception could be me saying something that misleads you. It could be me doing something that misleads you. So I'm facile enough to incorporate both of these things. Okay, the action taker, the receiver, is going to hear M, will not necessarily see X, so that's a distinction, and I want to make a distinction logically between X and M in the sense that the M, what I say, is something that isn't directly payoff relevant to you. There will be

an example on the next page that I think that will perhaps solidify that. The receiver takes an action, and they have payoffs, okay? In order to close this model, there's going to have to be a prior distribution on these unknown states.

Okay. This is a technical assumption. I'm not doing anything kind of technically fancy. Okay, what kind of model can you think of? You can think of a basic labor market signaling where the M is the observable amount of education I have. The theta is my marginal productivity, okay? In that model, I'm assuming basically that the education I have doesn't add to my productivity, doesn't change the -- my value to the labor market. So that's embedded in my framework.

Okay, cheap talk just means I see M, and it's not directly payoff-relevant to anybody. Verifiable information disclosure games are games in which I know something and I can say something about it and I can hide information, but I can't literally lie. So I have ten objects in my bag, and I can tell you I have an even number, I can tell you I have more than five, I can tell you I have less than 20, I can tell you I have ten, but I can't tell you I have 70. And then maybe more general games.

Okay, so I promised you a definition. What's the definition going to be? This. I'm talking to you. Mu. Mu is going to be how you interpret or decode what I say. So you end up -- you start -- come into this thinking maybe my product's good, maybe my product's bad. I say my product's good. You update, okay? It's how -- it's your decoder.

What's deception? Deception informally is me saying something that makes you think something that's wrong, okay? Now, there's going to be symbols here, and I got to put them up there, but the problem with that definition is that -- what does it mean for you to have wrong beliefs? So I want to argue that there's some subtleties here in that maybe we get a little bit of smart -- a little bit smarter thinking about those subtleties, okay?

So warning symbols, warning -- you can think of the two definitions that I write down as, I don't know, extreme cases. Okay, so this is one definition. A message is deceptive given what's true, theta, and what you believe, mu, if blah-blah-blah. What this is going to say is m is deceptive if I could have said n and you are somehow smarter if I say n. Mu is what you believe. That left-hand side is what you believe if I tell you m. I could have told you n, okay? That's on

the right-hand side. P is some number between 0 and 1.

What's that rho? Rho is bad. So the true state is theta. Rho, according to stuff that's written in those three lines, places all of its weight on things that are false. So the difference between what you believe if I tell you m and what you believe if I tell you n is that if I told you m, I'm taking the n belief and adding a bunch of stuff that's false, okay? So I'm misleading you in the sense that I could have told you something closer to the truth. Okay, that's one definition.

Here's another one. Strongly deceptive -- so this is harder -- is still m and n. M is what I say. N is what I could have said. Notice that the equations look kind of the same. There's a p in it and a 1 minus p and a mu and an n and a mu and an m. They're different, because the n is on the left-hand side instead of the right.

Here what's happening is I'm mixing m with that I business is with all the weight on the truth. So in the top line, number 1, I'm mixing beliefs with something that's bad. On the bottom line, I'm mixing beliefs with something that's good, okay? So if I say m and there's something else I could have said, n, this line says that n is closer to the truth.

Okay, if you don't like algebra, here are the pictures. This is picture number 1. So there are three possible states of the world, and the truth is up at the top. It's called theta, okay? And I could, if I wanted to, say something that will convince you that the three states are completely likely. That's the dot in the center. All those gray things, they are further from the truth, okay?

If I said anything that's gray, I'm deceptive in definition 1, okay? That's a definition. Again, what I want to get at is the idea that I'm telling you something that's further from something else that I could have told you that would be closer to the truth.

Okay, I gave you two definitions. You get two pictures. This is the second one. Before the belief that I pointed at was the n belief. Here, it's the m belief. If I tell you m, you think, oh, the states are equally likely, but anything on that line segment up to theta puts more weight on theta, more weight on the truth. So if I say m and I could have convinced you of something on that line, then I said something that was further away from the truth, okay?

In terms of the pictures, I want you to think that first one, there's a lot of stuff that can be deceptive. It's two-dimensional. The second one is

just that line segment, so the difference between strong deception and deception is that it's harder to be strongly deceptive.

Okay, there's a reason for this. I have some stories to tell. Okay, so one is this -- this may be my second joke, okay? -- so I think that's a commercial slogan, "Red Bull gives you wings." Now, I haven't defined lying, but I don't think when you drink out of the can white feathery things grow out of your back. So one could say that this statement is not true. It's a lie, but it's a metaphor, right?

I mean, you know, when you read it, you don't say, oh, I'm going to drink this and get wings, and deception is about what you believe. So if you read that line and you think Red Bull is an energy drink, and if you don't read the line, you think Red Bull is an energy drink, then this isn't deceptive. Pedantically, if no one believes your lies, then they're not deceptive.

This is, I think, a lead-in to my two important observations. I am going to repeat these observations when I come to damage, but let me say them now. Observation number one, according to my definition, I don't have to deceive you; that is to say, deception is about saying m when there's a better n. So I always

have the ability to say that n, okay? So if you're a particularly recalcitrant person, so here I am, I'm honest, I'm going to tell you everything about the product, and you think white is black, you get confused, I am not deceiving you, okay? I get inside your head, I figure out how you interpret my messages, and I tell you the one that's closest to the truth.

The second observation, you can be immune from being deceived. How? You refuse to believe me. So your mu is how you decode what I say. Suppose that whatever I say you just think that I haven't added any information to the picture. That means that mu dot given m is the same as u dot given n, so my definitions don't apply. So if you want to, you are immune from deception. If I want to, I'm immune from deceiving you.

You might say, all right, all right, all right, so these are simple properties, and I claim that they're pretty good desiderata for a definition of deception, okay? And it's important because this notion that you don't have to be deceived might have some weight, okay? I'm not going to argue this, but you could say, all right, well, then, it's your fault if you're deceived. Why should you be rejected?

With these two observations, think about

something else. If you protect yourself against all deception, you're also protecting yourself from all information. So it's not necessarily a good thing to put up this barrier, and just because I can avoid deceiving you doesn't mean that I'll want to avoid deceiving you. So if you leave yourself open to deception and our interests are different, then I might pull the wool over your eyes, okay? So I can avoid it, you can avoid it, but that doesn't mean that it will happen that way.

Okay, maybe a formal example will help. The reason that I give this formal example is to sort of emphasize the strategic flexibility I've got and to show you something that is counter some notions of deception. So this is a simple kind of disclosure game.

So the informed person is either going to learn 1, 2, 3, 4, 5, 6, 8, 9, 10, or she doesn't learn it, okay? So sometimes I'm informed about the state of the world, and sometimes I don't know. I get to say the state or I don't know. So your -- I'm going to approach you, and I'm either going to say "I don't know" or "theta is equal to 10," okay, but I can only say theta equals 10 when I know theta is equal to 10. Otherwise, you can take me to court, you can cut off my

fingers or something, okay?

So there are really 12 -- 11 things I can say, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, and I don't know. When I say 5, you know 5. When I say I don't know, you've got to think about it, okay? And you can think about it this way. You can say, well, gee, he'd tell me if he knew, so he doesn't know, so I'll sort of believe the prior; or you could say he's strategic, and if he knew it was 10 and 10 was really good for him, he'd tell me that.

So in standard models, I like you to believe that theta is high, then you like to know the truth, and you might imagine an equilibrium in the following form. If I see 10, I'll tell you. If I see 9, I'll tell you. If I see 8, I'll tell you. If I see 3, I'll say I don't know, okay? So why do I do that? Because if I say I don't know when it's 3, even in equilibrium, you'll scratch your head and you'll say, is it low or did he really not know? And maybe you'll think that it's sort of an average between a low thing and 5, and I'll be better off with that average even if I -- even in some cases, okay?

So in this case -- I sort of hinted at these preferences. In this case, I'm going to say I don't know when the type is small, zero -- sorry, 1 and 2

maybe, 1, 2, and 3. What are you going believe? In equilibrium, if I say 10, you believe 10. If I say 9, you believe 9. If I say 8, you believe 8. If I say 2, you'll believe 2, but in equilibrium, I'll never say 2. If I say I don't know, you'll believe that it's an average between 5 or 5 ½, what you think, if I really didn't know, and 1 or 2, because you're sophisticated, you know, sometimes I'll say I don't know when I really do know, okay?

When I say I don't know in that situation, I'm being deceptive, okay, because I can push you closer to the truth by actually revealing, all right? Why do I do that? I wanted to illustrate the notion of deception, and I wanted to illustrate that this notion happens even in an equilibrium, okay? So everybody's sophisticated, but I'm deceiving you because I could have said something that was better, more informative, and I refused.

Okay. Silence and omission, so I think that a lot of people view silence as salient, and they say, well, you know, we didn't say anything, so he couldn't have deceived me, and I want to argue that that's a trap, and it's a trap in theory certainly and in practice maybe.

So the theoretical trap is the following:

First, there's no such thing as omission in the formal model, so, you know, it's -- you make it up. I just -- it's an m. I didn't say m was omission. I didn't say m was equal to theta. Okay, that's no problem; that is to say, in the example that we just saw, "I don't know" could have been interpreted as silence.

You know, I come to you, and you wait to hear how many things I've got, and I either say five or I say -- and you can say, okay, well, either he's silent because he doesn't know anything or he's silent because he's suppressing the fact that he knows something bad. So there's a sense in which I can add silence to the model.

But now whether this is deceptive or not is going to depend largely on how you interpret it, okay? And you don't have to interpret my silence as the prior. You don't have to interpret my silence as something that's absolutely revealing. And whether this is deceptive or not depends. So modify the previous example, make it easier.

Suppose that you're sure that I know what theta is. In that case, the standard equilibrium, the equilibrium that passes refinements, is one where you're skeptical; that is to say, if I don't tell you the state, you assume the worst, okay? In that case, I

can't deceive you, okay? You interpret silence as, oh, theta is equal to 1, and you move on.

So what I'm saying here first as a sort of theoretical model, you say silence -- silence doesn't mean anything special in the formal model. And then I say, okay, I'm open-minded. I can make the model such that silence or omission means something, and then silence can or need not be deceptive for the same reasons as any other message can or need not be deceptive, all right? And maybe there's something here. Okay.

So it's the beliefs that determine whether a distinguished message, silence, is deceptive or not, and that's something that sort of operates at a different level than what you call it.

Okay. I didn't search this. I made it up. So I hope there's no such real thing as Zyllyz, but let's think about this as a way of thinking about what's going on here. So I can make several announcements. Announcement number one, I can say, "Buy Zyllyz." Announcement number two, I can say, "Zyllyz helps you think forward and backward." Announcement number three, I can say, "Eight out of ten economic theorists can become industrial organization economists with regular treatment with Zyllyz." So I haven't tried it,

sorry, but -- so there's a sense in which I just want to put this out to sort of interpret what's happening in my land, and the answer is a sort of "I don't know."

So I'd say that "Buy Zyllyz" is basically silent, okay? It's not intuitively dishonest, and it could be deceptive, but being deceptive requires sort of assumptions about the cognitive process that's going on. So when is it deceptive? It's deceptive if I know that I've got a really bad product, and when I say "Buy Zyllyz," you either think, oh, it's worth buying, or there's no particular information in that. You're basing it on a prior.

Okay. So conceivably, just a statement to buy is deceptive, and this shows sort of how much weight is placed on this mu, and I hope that before I end -- I'm not sure I'll get there -- but I hope that I'll say something about, you know, that's the practical consideration.

Again, there's no operational meaning that I know of. I don't know what thinking forward and backward means. That's basically the first one. And the third one could actually be demonstrably false, but, again, whether it's deceptive or not is based on how you think about it.

Okay. Up until now, I've talked about

deception as being about messing with what you think, but we ought to think about deception in the context of who's helped or hurt by these statements, and that's the notion of damage. So I've got one definition of damage, and there's symbols here, but basically, with deception, m was deceptive if I could have said n, and when I said n, you had better beliefs. M is deceptive if I could have said n, and when I say n, you get higher utility, okay? So I expressed that in notation. This step has less ambiguity than the definition of deception. It's more straightforward, maybe more definitions.

So here I'm talking about both m and x. You can forget the x. The U-upper-bar-r is how much utility you get. I say something, m. You take a response, and then you get some utility. Okay, m is damaging to you if I could have said something else and you'd be better, okay?

So proposition. This is a version of the proposition about deception. As with deception, you don't have to be damaged, and I don't have to damage you, okay? If you ignore what I tell you, then you won't be damaged. If I pick the message that leads to your highest utility, then I don't have to damage you. But, again, if you ignore what I tell you, then you

pass up opportunities to learn from me, and if I don't damage you, then I might be damaging myself. So just because those things are avoidable doesn't mean that we will avoid them. So those propositions are, I think, canonical, they're simple, and they're useful desiderata for the concept.

So you got your damage. You got your deception. The same guy is giving the talk. Probably they're related. And there are mathematical propositions here, and the mathematical propositions are basically of this form. If I deceive you, I'm damaging you, and if I damage you, then there is some deception involved.

Now, I gave you two definitions of deception, so it has to be the case that there's somehow a way in which you can modify the informal statement in quotes to fit both suits, okay? And I claim that that's possible, and that's a technicality, and, you know, I'm here tomorrow. I'll tell you about it.

Okay, the difference between deception and strong deception. Deception is easier -- remember the picture with the triangle shaded and the picture with a line shaded? -- and the specification of preferences can come from a smaller class, and that's the thing that I skipped, but basically when I deceive, I damage

you for all preferences in a class, and when I change the definition from deception to strong deception, I change that class.

Okay, other things you can talk about. There are different ways of talking about "further from the truth." I think I gave you two polar cases. Other cases come basically if you know more information about the structure of states. So if it turns out that state one is close to state two, then I can come up with a different notion of deceiving.

Okay, prior beliefs and optimal actions as a benchmark. There are lots of words here, but I think that there's a strong intuition that maybe I'm deceiving you if what I tell you leads to something that's bad relative to the prior information. I want to argue that that's wrong, and it's wrong because those two simple properties that I gave you won't be true anymore, okay?

So if I'm going to be judged on whether I deceive you relative to the prior information, well, it could be that there's nothing that I can do that will leave you in the same situation as the prior, okay? So I won't be able to avoid deceiving or damaging you.

Finally, really, the standard definition from FTC of deception includes the spirit of inducing wrong

beliefs. They use the term "misleading." It includes damage, but I've connected damage to deception, so I'm really somewhat in line, and the one difference is that, quite reasonably, the natural definition from the FTC tells you something about where the mu comes from. So the beliefs of statements somehow are nailed down by some notion of what is reasonable, okay? And for a game theorist, they're just there, okay?

So this will be my end page. I hope -- I think I've gone too far, but I'm within two minutes of the time limit. So one question is, where do these beliefs come from? And for theory, anywhere, I get to make them up. For practice, you want to do something like find somebody who's representative. So one notion is that people take statements literally. And the Mini Wheat case is a case in which experimental tests were framed in a way that was misleading, and so benchmark beliefs would be just to take messages literally.

I say something about this, but let me skip it. I think that I -- these are interesting things to leave on the screen, and if nobody has questions, I can tell you what they mean, but I've gone too far. So I will say that's the end. Thank you.

(Applause.)

MR. SOBEL: So one of my advantages, right, is

that I don't have a discussant, but you can --

MR. THOMAS: We have about five minutes for questions.

AUDIENCE: I have a question about the -- first of all, it was a really interesting talk, I thought, but I have a question about this representative consumer, because what I was thinking before you put that slide up was if deception is about how the receiver interprets -- which is a perfectly reasonable way of looking at it -- and if receivers -- consumers, if you want -- are heterogenous, which undoubtedly they are in terms of their cognitive processing of whatever they are -- let's just call that intelligence for the moment -- then, A, it's not clear to me if there is a representative consumer.

You might have bimodality, people who are very intelligent, people who are not. Even if you did have, you know, say a normal distribution, you know, the question is, what are the social costs of deceiving those people at the low end of the distribution? That might be very considerable, and yet if you base it on a representative consumer, it's -- or receiver of the message, we might say that it's not deceptive.

So have you thought about that aspect, kind of the welfare aspect here?

MR. SOBEL: So I don't think I can directly answer the welfare question. The dodge is the following: The policy issue is what's your benchmark for deceptive statements? So the issue really is -- let's take this disclosure game, okay? Sophisticated folks have read Grossman and Milgrom and Roberts, and they understand about skepticism. So they'll know that if a statement is of the form, if there's, you know, more than $5 in this bag, then they should think 6.

So if I am going to protect consumers, should I have a standard that is flexible enough to include folks who take that statement literally, that is to say, well, it's more than 5, it can be anything that's more than 5, so I will average 6, 7, 8, 9, 10? And that's a policy issue.

The cost of not protecting depends, of course, on how many people aren't skeptical and what happens to those people for being not skeptical. The benefits are, you know, how you repair those folks, and here I'm hiding; that is to say, as a theorist, I can say, you've got to think about that stuff, and I can't tell you what's right. It depends on who you want to protect and how many of them there are.

AUDIENCE: (Off mic) -- maybe I should have stated it this way before. So there's the welfare

aspect. That's fine, I take your point. But if I say something -- if I give you a message m, and I know that there's a distribution of consumers, and I know that only a few of them will misinterpret it, then I -- it would seem to me that I would want to call that deception, because I don't know who you are, but I know that you're out there, and so somebody will, in fact, be deceived, so that the sender of the message, if he knows that there are heterogenous consumers in terms of their cognitive interpretation of this, it seems that if mass of those consumers, that that's deception.

MR. SOBEL: So I listened to your question. My reaction to that is going to be a stuffy, formal one, and it is I would call that statement deceptive only if -- but if and only if -- there was another way that I could reach those naive consumers and get them closer to the truth, okay?

So if I could only make one statement to everybody and it turns out that that one statement is going to be misinterpreted by a fraction of the population no matter what, then I'm not going to necessarily call it deceptive. If it's the case in this disclosure game that if I, you know, literally said the number -- rather than greater than 5, I said 6 -- and I get everybody on the right page, then I'm

inclined to call the vague statement deceptive in a world where there are naive or unsophisticated folks.

You're on.

AUDIENCE: Good. I was just a little bit concerned on some of the language you're using because the sender is being the one who damages in the sense you can say m or n, and that can be damaging, and it can be damaging to the receiver or to him if he chooses something that's bad for him, but communication is both ways. So you -- so because I understand deception, I understand damage, but my concern is with the blame, because you're always blaming the sender, whereas communication is both ways. So you need a message and you need an interpretation.

So me, as a sender who would like happiness to be around, maybe I would like the receiver to change his interpretation when I cannot do that. I wouldn't like to be blamed because of how he is interpreting, like, my --

MR. SOBEL: So my sender knows how her audience is going to respond. So it could very well be that her audience is silly, is -- you know, doesn't take good actions with available information, but the communicator can trace through the process and can figure out the implications of her words.

Damage happens when she traces through that process and still does something that's bad for the audience, okay? So it's not going both ways in two senses. One, the communicator knows how the information's going to be used, and she can adjust for it; and two, we're not thinking about the sender's preferences at all.

For my next talk, I can talk about how deception influences the sender's preferences, and there the jargon words are charades and bluffs. So invite me back next year and I'll -- okay, you're on.

(Applause.)

MR. THOMAS: All right, yes. Thank you very much, Joel.

Now we'll break for lunch, which is sponsored by the Tobin Center for Economic Policy, and there is food in the lobby. Let's reconvene at 12:45 for a keynote address by Panle Jia Barwick. Thanks.

(Whereupon, a lunch recess was taken.)

AFTERNOON SESSION

MR. VIOLETTE: Okay. Hi, everyone. Welcome back from lunch. We're excited to have Dr. Panle Jia Barwick presenting for the keynote, and she is an Associate Professor of Industrial Organization, Applied Econometrics, and Applied Microeconomics, in the Department of Economics at Cornell University. She serves as the co-director of the Cornell Institute for China Economic Research.

She's also a faculty research associate at the National Bureau of Economic Research, an editorial board member of VoxChina, and an associate editor of the China Economic Review. She received her Ph.D. in economics from Yale University, and her keynote today is titled, "China's Industrial Policy and Empirical Evaluation." Thank you.

MS. BARWICK: Thank you, Will, for your kind introduction, and this is -- also, thanks to the organizers for this great opportunity to talk about our work on China's industrial policy. This paper is one of several projects that I have -- industrial policies, broadly defined, including -- so here is about subsidies. Other papers include market exchange with technology and other industry-based policies. So the paper here is joint with Myrto Kalouptsidi at Harvard

and Nahim Zahur, who is a Ph.D. student from Cornell who is on the market this year.

So this is probably a very familiar picture to many of you here, where China experienced exponential growth in many factory industries in the past several decades. What is remarkable is the speed at which China's takeover is happening. So, for example, just to give you a few examples here, China's world market share in solar panels increased from 5 percent to 55 percent in five years, changed from a net importer of steel to the largest exporter in three years, doubled its world market share of shipbuilding in two years, and overtook U.S. to be the largest auto producer in 2009. So this is -- here's a list of, you know, several, like, industries where all of them grew by 20 to 30 times over a 15-year period.

So there are many reasons that explain the exponential growth, including the access to WTO, the various reforms, but we also argue in our paper that some of the expansions are fueled by the massive industrial policy in China. The famous ones that many of you probably have heard of include the national and regional five-year plans, where certain industries are designated as a strategic or pillar industry that receive preferential treatment from the government, and

the recent one, which is also a bit controversial, is Made in China in 2025, where the government wants to dominate in the ten industries of the future, including artificial intelligence, aerospace technology, and clean energy cars, as well as the latest generation marine ships and marine equipment, so that we also have several papers on clean energy cars, where the policy has been -- the magnitude of policy has just been massive.

And as a consequence of those policies, you also see a pattern where the industries typically have low concentration compared to international standard, access capacity, et cetera, that we will also -- you will also see this pattern here in this study.

Industrial policies are actually pretty -- quite prevalent, and famous examples include U.S. and Europe after the World War II; Japan, South Korea, and Taiwan from the fifties to the eighties; China, India, Brazil and many other developing countries in the last couple of decades.

Despite the prevalence of industrial policies, there aren't many empirical research that study the welfare implications of industrial policies. A lot of the studies describe what is happening to the tired industries in terms of output and revenue, but in terms

of the cost and benefit analysis using microlevel data, there's actually not many and, you know, despite the very heated debate about industrial policies, and that's where this paper comes in.

So we are going to use data about the global shipbuilding industry. We collect data on all of the major shipyards, on their quantity and price and products, and use shipbuilding as a case study to illustrate -- first, we are going to do a quantification exercise to show you what are the magnitude and the cost and benefit of this policy.

And then, more importantly, we want to think about general guidance, you know, some of the lessons we can learn or implementation, you know, if you think those policies are desirable, how do you implement it? That's the goal of our analysis.

So that requires counterfactual analysis where we simulate different combinations of the policy mix, and just to give you some summary findings, first, as I mentioned, the magnitude of the policy relative to the size of the industry is very big, and you can see that the aggregate industry revenue is about 1700 billion RMB, and the numbers I put there is subsidies for production, investment, and entry, total about half a -- half a trillion RMB over our sample period.

And the production suddenly is about 159 billion RMB, followed by investment, 51 billion RMB, and then the largest amount is entry, which is 330 billion RMB, mostly cheap, vastly subsidized land that were given to firms. Those policies boosted China's investment.

By the way, the timer is not right.

MR. VIOLETTE: We will give you the time.

MS. BARWICK: Okay.

So it boosted China's invest -- I have still 30 minutes, okay, thank you -- by 270 percent and entry by 200 percent. It enhanced China's world market share by 40 percent. It has a huge impact on world ship prices.

Now, surprisingly, we are basically going to decompose the impact on different products, and more importantly, we want to see which policies are, you know, more effective or create less distortions, and it turned out that the production investment subsidies, they have different tradeoffs, but they can be largely justified based on the revenue or output considerations.

Well, entry subsidies are extremely wasteful. So I will show the numbers. And the industry profit -- you know, I show you a big number on the subsidy, but industry profit in the long run only increased by $145

RMB, so the net return is actually quite low, and the subsidy reduces concentration because of this flood of entry and lowers capacity utilization, but the insights that we want to mention, to highlight, are three of them, three key insights.

First, distortions are convex. By this what I mean is not only the distortion increases convexly when the magnitude increases, but also, when multiple policies interact, the impact is worse.

And second, it actually strike us by surprise the magnitude. It turns out that dynamic sorting is extremely influential, important, the same policy, and then the cyclicality is highly relevant. I will show you the numbers. The return can differ by three times if you just change the timing of the implementation, and I'll explain why that's the case.

Okay. So I won't finish the paper, it's a very long paper, so I probably will skip many of the details. I will give -- introduce a little bit of -- describe the industry a little bit and skipping estimation entirely, although that takes years of hard work, and talk a little bit about counterfactual and the lessons that we're learning.

Okay, so this is global shipbuilding, but it's mostly about China, because that's where the action

takes place. The other countries, the production in the other countries are more or less similar, where China experienced a massive expansion during this period.

So this is a graph about the market share, essentially, the percentage of ships launched over the past hundred years or so. So shipbuilding is a classic target for industrial policies, partly because of the commercial and the military considerations and throughout the last century. So for the -- a big chunk of the last century, Europe is a major producer, followed by Japan, South Korea, and now we see China.

So here's -- so in 2005, China designated shipbuilding as a strategic industry, a pillar industry, and the scope and frequency of national policies starting from 2005 is unprecedented. So I have listed a whole -- you know, a large number of policies here that basically target both shipbuilding and ship repair, ship equipment, and all of the facilitating sectors. And many of those policies have specific output and investment targets that the industry should meet.

Quickly, the industry ballooned, and there was signs of overheating. So then in 2009, the government basically put a ban on entry -- no new projects would

be approved -- while strengthening support for the existing firms. So those are the two, the 2009 -- the implementation of the major subsidies, and then 2009, when entry was banned, are the basic two shocks that we focused on in this analysis, where our sample period spans from 1998 to 2013.

So this is the chart about China's increasing market share in this industry, which was around 10 percent in the beginning of our sample and quickly, as you can see, took over to become the largest producer of ships by output.

And I want to emphasize that for a big chunk of the sample period, Japan and South Korea actually were larger producers, but as you can see, there was minimal entry and minimal expansion in these other two countries. On average, the entry of new shipyards in Japan and South Korea is about 1.4 per year, where in China, the number of new shipyards -- those are large shipyards -- producing ships for sea navigation exceeded 40 -- 30 or 40 per year. So there was actually massive expansion for this industry.

And there's no other picture that's more telling of the magnitude of the -- you know, the expansion that the industry -- aggregate industry investment increased by four times nearly overnight in

2006, and there's also, you know, very strong evidence that the expansion is driven by government policy rather than economic considerations, because there are very few economic arguments you can provide that can justify this kind of picture. The expansion happened across the board, not just at a particular region, but across all kinds of ownership status, all young and old firms, and across many provinces.

Okay. So the model is a reasonably standard dynamic model where we spent a lot of time improving or at least extending the existing literature on investment, because that's where the action takes place, but in the interest of time, I won't get into the details at all. It's a very complex model, but I'll basically describe what we do and then leave the rest to -- you can read the paper, which is, you know, quite long.

So we have all of the firms. We have Chinese firms and firms in Japan and South Korea, and the other firms in China make dynamic decisions, meaning they decide whether to enter, and they decide how much investment to undertake, and they decide whether to exit or not. That's largely because of data.

As I already mentioned -- explained, there's very little expansion in Japan and South Korea. So

firms in South Korea and Japan, they make production decisions, which are static, but they are not making dynamic decisions.

Okay, and then we have -- so we have dynamic decisions in investment, entry and exit, static decisions in production, and we have different ship types. So here we focus on three, which are bulk, tankers, and containers, which are essentially 90 or 95 percent of the total ships ordered.

And then those three countries explain also 90 to 95 percent of world production. So we're actually focusing, indeed, on the world industry. So there are three -- broadly speaking, three kinds of policies. The first one, we call them production subsidies. These are essentially -- you can think of subsidized input, which is quite standard, and export credits, because these are ships that are trading in the export market, and there's a very important word, which is a preferential buyer financing.

It used to be the case that shipyards, in order to attract buyers, would provide financial services, because ships are huge commodities, you know, they are large products, but during the sample period, China set up banks to provide collateral loans on behalf of the firms at very favorable rates, essentially heavily

subsidized rates. So that we treat as production subsidy because it lowers the cost of servicing customers.

Then the capital subsidies are -- you know, you probably have heard of. Those are essentially low-interest credit, long-term credits, and a very important one is the tax credits for accelerated depreciation. So those are actually quite important incentives for the firms.

Finally, entry subsidies are cheap land, you know, essentially free land given to a large number of shipyards simply by registration procedures.

And then in the model we have a fairly simple model of the timing. So in 2006, you can see the policy kicked in, and then in 2009, the entry subsidies were removed, but the other two remained. So we allow the magnitudes to differ across different sample periods.

Okay. So then the static -- so basically you first use their production decisions to back out the cost of production. How do we do that? We basically look at the variation in price. So here if you said marginal revenue equal to marginal cost, and we know the marginal revenue because we see the fluctuation in prices, and we have a reasonable sense of demand, then

that can trace out the curve of marginal cost, okay? So that's where we do -- estimation of the marginal cost of production.

And then for the dynamic decisions, we basically look at firms -- this is a little harder. So firms are going to make optimal investment decisions to maximize their lifetime profit, okay, lifetime revenue. How do -- the key of our exercise is to figure out what is the actual investment cost firms face, and this is actually where part of our innovation comes in, where we focus on the macro literature but also extended to account for heterogeneity among firms.

I won't go into details since I only have 30 minutes left, okay, so you trust me that I did a good job at estimating, you know, the cost of investment, and skip the data and skip the estimation.

Okay, so here's the production cost. We allow the production cost to differ for each ship type, which is very important. The (indiscernible) could be different. Then we did find that the marginal cost is upward sloping, and -- but then the magnitude of the subsidy is about 10 to 13 percent of the price. So think of that's the amount of subsidy the firms receive from their production.

And then we actually did spend a lot of time

doing robustness, thinking about there's -- one key argument of industrial policy is that it has to be some kind of externality, so that's where we did -- you know, spent quite some time. We find very limited evidence -- even on the most conservative side, very limited evidence of some kind of -- any kind of spillover, okay?

And then for the -- then we estimate the investment basically, again, looking at the actual investment firms are taking, tracing it against the potential increase in their revenue, that's how we estimate the cost -- the investment cost, and we also did a quite job.

So going back, you can see that the magnitude literally is if you invest a dollar, 25 cents will be subsidized, and later on, post-2009, the magnitude is even higher. The reason is that the financial -- after financial crisis, the ship prices were plummeting, yet you -- but still, the firms are still expanding, investing, and that's a sign of strengthened support from the government. So our feeling is that is pretty -- and we have the entry cost estimates, which is about $2 billion RMB. This is similar to the accounting measures, as well as our scrap value estimate, quite similar to some of the merger and

acquisition numbers that we got.

So then with this we find out -- we also estimate demand. Demand is useful for counterfactual, because in a counterfactual, we are going to intersect global demand with global supply, and that's where we're going to simulate what happens to this industry.

Okay, so to evaluate the industry, it's actually quite important to take the dynamic consequences into account, because both entry and investment have dynamic consequences. Today's expanded capacity will influence tomorrow's price and firm activity. So that's why in order to evaluate the industrial policy, it's not sufficient just to look at what happens that year; rather, look at what happens the next 10 or 20 years when the capacity or when the accumulated capacity is still productive.

So what we do is we show you a bunch of different simulations that basically turn on and off different subsidies and let the industry evolve for a long period of type. So here we allowed the industry to evolve for a hundred years, but you can actually just look at what happens for the first ten years, what happens for the first 20 years. The qualitative results look very similar, okay?

So this is just to show you, if we turn off the

subsidies, what would happen. You can see that entry would still occur because China has many undeveloped coastal lands, and the ship prices were quite lucrative but not nearly as much. Similarly, for investment, investment, indeed, will expand, but as you can see, once they elevated, when the ship prices plummet, and so then this is also impact on concentration.

So here's the impact on world ship prices that I mentioned. So without a subsidy, the ship prices will be about 40 to 70 percent higher in the absence of a subsidy, okay, which is a big supply expansion. And then so it increased China's world market share by 40 percent and benefit ship -- worldwide shippers -- this is actually quite a big number -- $230 billion maybe, but China domestic firms only benefited about 10 percent, because China is quite a small player in global shipping.

And then we -- so here's the interesting part. We want to understand what are the relative accuracy of different policies. So what we do is we turn on each one of the policies at a time. We first turn on production subsidies, several investment sub -- several entries, see what happens.

Okay, so the very last column is when there's no subsidy, nothing. The industry is just evolving on

its own. The first column is when we have all subsidies. There's several important messages. The first, if you look at -- so these last three in purple, the very bottom one is the net producer surplus. Here is the revenue minus variable production cost minus entry cost minus investment cost, plus scrap value. So the very last bottom one is the one you should look at.

You can see that investment and production don't do terribly bad. If you look at the -- this row here, this is actually revenue. So the increase in industry revenue, lifetime industry revenue, as a ratio of the dollar of subsidy you put in, if you use output as a benchmark, both investment and production subsidies actually are doing okay.

And second is if you look at the entry subsidy, you can see that the return is very low. The reason is it attracted a large number of firms that, if you look, are not very good. Indeed, those are the firms that were first to exit after the industry collapsed.

And then if you look at the -- all of the returns, the return for all subsidies, that's much lower than any one in isolation. That's what I meant by convexity, where there is a more thorough analysis of that. So then the subsidies were much more distortionary when combined.

The second one is we also compared the production versus investment. In the interest of time, I will basically just quickly summarize. The production, not surprising, is more effective by raising revenue in the long term, over the long period of time. Investment tends to target better firms that are more efficient which, as a result, has a long return -- you know, has a higher return over long period of time. The reason is that the better firms are likely -- are more likely to invest, while the firms that are not very efficient, they are more likely to exit and they are less likely to invest.

This is actually -- so this is a decomposition, where we look at the amount of subsidies that were taken out by firms and do a very simple decomposition by looking at how much subsidies were taken out by firms that were efficient versus inefficient. Those efficiencies are just looking at observables, looking at a firm's observable attributes that are correlated with their cost of production if you are efficient firm, meaning your marginal cost of production is below average, okay?

So you can see that across the three given subsidies, production, investment, and entry, entry is equally taken up by good and bad firms, while

production and investment are much more likely to be taken up by firms that are efficient, and that's actually a key reason that explains the differences in this return. And as a result, if you -- if you can -- so that's what explains the relative efficacy of production investment. Investments, in particular, are much more favorable. Only firms who are productive, more likely to stay active in downturn, will take up investment subsidies.

And this is actually a picture -- a table that strike -- that took us by surprise. We didn't anticipate this result. What we did, you can see the two simulations. The first one is we -- both of them have the exact same amount of subsidies. The only difference is one you hand out the subsidies. The first one is we gave out all the subsidies given the boom. The second one is we gave out all the subsidies during the recession, after the two thousand -- after 2009.

You can see the bottom line, the percentage -- the bang for the buck, if you want -- if you like, differs by three times. So putting a dollar -- if you give firms a dollar in the downtime, you get about 78 cents in return in the lifetime of profit, while the first one, there's only 29 cents. So what explains

this very different magnitude? There are about four reasons.

The first is that the composition of firms in those two periods actually are different. The firms that are active, that are both producing and less likely to exit during downturn, are the good ones, and a that's the first one.

And second, expansion is actually quite costly during the boom. During the booming period, firms are already producing at capacity. If you want them to expand even further, you are only driving up the -- you know, the rising part of marginal cost.

And the third that I didn't actually mention here is that there are scales of economy, and the active -- the good firms that are active, they're efficient, they benefit more from scale economy when -- during the downturn period.

So the actual policy, if you observe -- if you look at the data, it's highly procyclical, and that's why the impact is -- you know, is actually that greatly influenced, reduce the effectiveness of the policy. And to be fair, having countercyclical policies are hard, because, you know, the government doesn't have enough revenue, but in this case China actually have five -- has five-year plans, where they can -- they

have a lot to do in the five-year and ten-year plans, where they can balance the cyclicality of the industries.

And so just to show you the differences in firm composition, this picture basically plot the efficiency, which is essentially the part of firm marginal cost separated by -- for the two counterfactuals. The blue one is the average firm efficiency, the average of the firm efficiency for if you -- if you subsidize firms during the downturn. You can see that actually they are much more efficient than otherwise just because the bad firms aren't active during the booming period -- during the downturn.

And finally, so we added in quite a bit more in the paper. We also look at the consolidation policies. We also evaluate different combinations, and I will leave it to you to read the paper. We spent a lot of time trying to think about rationales, why would government do this in light of some of the negative findings we have, and so first we looked for the traditional justification for industrial policies, which is externalities, how -- but we find that shipbuilding, particularly for bulk and tankers, those are fairly mature technologies. You actually -- we don't find any evidence of industrial -- you know, the

learning by doing or spillover effect.

There is -- what is inconclusive is for containers, where, indeed, there might be some learning, but -- that's the relatively high-valued product, but just because of the limited observations, you know, we did, there aren't many containers produced in sample period. We couldn't really say much. But more concrete evidence for bulk and tankers, which also absorbed the lion's share of the subsidies, you know, there is no evidence of spillover.

The second one, which is often cited in the literature, is a strategic trade consideration. You're basically competing with other countries. You want to subsidize your firm so that you can grab more -- much -- a bigger chunk of the rent. It turned out, for strategic trade policies to be effective, the industry has to be grossly concentrated. There has to be enough rent on the table for you to grab.

In other words, you know, another example would be the high-speed rail that's, indeed, where the only few players in subsidizing the capacity of firms have a big impact on the rivals. This industry has more than 400 players. Even the largest counts for less than 5 percent of world market share. It's -- there's just no rent, available rent for strategic trade policies to be

effective.

And, third, spillover to other sectors. There still is a very small -- sorry, shipbuilding is a very small downstream for steel -- for the steel industry and also counts, you know, less than 5 percent -- 0.5 percent of national employment, so there's very little spillover to other industries, and China is also not a big player in international transportation services. That leaves open two other potential justifications.

One is the impact on trade. Indeed, subsidized shipbuilding lowers ship prices which then can lower freight rate which then can boost export and import. The reason we stop there is because the literature doesn't give us a number on the welfare benefits of expanding your trade, and so that's where we basically -- I think there might be some benefits that we can't quantify.

And the last one I think is the military and, you know, national security considerations. Whatever the motivations are, we provide you the cost analysis to show you here's the price you pay for achieving the output that you want to achieve.

Okay, just to conclude, the broader license we have is that the magnitude of policies is massive and that, more importantly, the effectiveness of different

policies is mixed. We want -- so the production investment subsidies can be justified by output considerations, but entry subsidies are very wasteful. That's, I think, why partly the Government just changed it, you know, latter part of the sample, and then the production study and investment study, they have their tradeoffs depending on the policy goal.

Finally, the distortions are highly convex and deteriorate with the magnitude, and also the number of subsidies in place, the dynamic sorting and targeting is very instrumental. And finally, in many factory entries, the cyclicality plays a very important role, and that affects greatly the advocacy of the policy.

Thanks. I'm over.

(Applause.)

MS. BARWICK: Okay, I think there is time for a few questions. Please.

AUDIENCE: Hi, Steven Bristoll, Federal Maritime Commission. Actually, China is one of the biggest consumers in at least container ships. They -- the Chinese Government owns one of the biggest players, Costco. So wouldn't they want to have an incentive, say, for them to buy their own ships, and couldn't that explain a lot of this aggressive expansion?

I know the Chinese Government is also doing

things like buying up ports in other countries, rapidly trying to expand their territorial waters. Couldn't this all be part of that sort of game that they're playing?

MS. BARWICK: Yeah. So we looked into the freight service, which is what you're alluding to. If you look at, in the grand scheme of things, China accounted -- used to account only 3 percent, and now it account around 10 percent of the total freight service. So that's why a big chunk of the benefit of this policy actually goes to international foreign firms, not Chinese firms.

That's where we showed that the world shippers benefit by about 230 billion MNB, but domestic firms benefit about 10 percent of that. That's, I think, where the number -- how the number comes -- where the number comes from.

And second, you know, we don't want to -- it's -- say that the policy have no other implications. You know, I already mentioned several, the military and national security. There could be important considerations in trade, too. The reason we stopped to say -- to give you the dollar number on the welfare benefit of expanding your trade in the -- doesn't have a conclusive number.

AUDIENCE: Nobody knows.

MS. BARWICK: Right. So what we want to show is, regardless of your objective, if you think industrial policies are desirable for whatever reason, implementation is very important, and we basically provide some general lessons on how you can make it more effective. Yeah.

AUDIENCE: Okay. Thank you.

MS. BARWICK: Thanks.

AUDIENCE: Hi, Panle. Can I ask, like -- it's kind of a crazy question, but -- so back a long time ago, Michael Boskin got into trouble for saying, like, it doesn't matter whether we invest in computer chips or potato chips. I took this as a sign that I should do research on potato chips, but -- you know, but most people have taken that to mean, like, industrial policy is going to be largely irrelevant, at least in the U.S., and we shouldn't bother with it.

And so it seems like you're saying still sort of less than -- you were still at cents at a dollar, so it still seems like in net it's not, but you're saying we could make it less bad. Is that -- but should we be thinking about this at all or not, I guess?

MS. BARWICK: Right. So there are many different kinds of -- so industrial policy is very

broad. For example, one kind is a place-based policy, is actually very active, and in terms of targeting industries, so traditionally the literature actually has quite a few justifications, most of them about externalities. So what should you do? Obviously you should target industries that have a huge spillover effect.

So that's why I don't want to say -- use this as example to say, look, industrial policies are totally bad, but that's why I think that, depending on policy objectives, you might want to do it for other reasons, and we point you to directions of how you can do it better. Yeah.

And actually, there are -- even in U.S., there are other industries that are -- obviously you want to favor domestic producers, you know, depending on objectives, and you can design the different policies such that it benefit -- they benefit domestic producers.

AUDIENCE: Hi. Thank you for your presentation.

Just a quick note on the effect of policies that encourage entry, and you cited examples of such a policy as providing free land, for example. Given that ownership rights and property rights in China is still

a fluid area and there may be instances where property rights are not well defined, could that explain the ineffective -- the relative ineffectiveness of entry based on the provision of, you know, free land?

MS. BARWICK: I think that's a very important question. Well, where the -- actually, the implementation of the property rights is actually quite complex and differs greatly across different provinces. What we're saying, the reason that entry subsidies are ineffective is because the -- you basically are giving out subsidies to everybody, including firms that don't have a track record -- including firms that really haven't, you know, proven themselves. It's not selected yet, you know?

That's a stage where everybody can get subsidy, but for production investment, you can see that the good firms are actually more active in taking up these policies. So that's a lesson we have learned. If you can design a policy such that the take-up rate is higher among firms that are more efficient, that will make the policy better, and I think that's another reason -- so here, a lot of -- as I mentioned in the talk, a lot of the firms entered because it was a lucrative market to be in in 2005 and '6, but quickly, after the -- when the ship prices plummeted, those

firms that entered because of subsidy were the first to go out.

AUDIENCE: Thank you very much for a very interesting presentation. I have two questions. One is, these kinds of subsidies, of course, should be counteracted under the World Trade Organization, and I know there's been discussion of that. Are there any shipbuilding cases from Japan or Korea against the Chinese, and if so, could you bring me up to date on those?

The second point I really found interesting is that subsidies to new entrants were so -- were so less efficient, so less effective, and I was wondering if the new entrants were often in some of the weaker provinces, because you noted at the beginning that you have regional policies as well, and one of the things I found in my work in China over the last 30 years is that there's a tremendous difference in provincial capabilities in China.

So I'm wondering if the new entrants would have been more clustered in some weaker provinces that had less capabilities and so on and were maybe the wrong places for those entrants to be going in. Thank you very much.

MS. BARWICK: Thanks. So we -- you know,

basically, it would be a very important question to look at the performance across different regions, and we tried to look into that, but the data is not rich enough for us to say are those policies -- first, to quantify the different magnitude that went into different provinces, and second, to say -- to evaluate the effectiveness.

So that's where I think we stopped. We haven't pursued it enough to figure out -- actually, largely because of data considerations, that we couldn't really give you a precise answer about the extent of the subsidy magnitude in different provinces. So that's actually our -- you know, I just want to acknowledge that.

What's your second part? I forgot.

AUDIENCE: Subsidy (off mic) WTO.

MS. BARWICK: Oh, indeed, there are cases. I don't want to say something I don't know for sure, so I will probably do some research and tell you later. I am aware it's quite -- you know, quite -- part of that is confidential, but to the extent I can, I will follow up with you later.

AUDIENCE: Panle, I have a couple of -- first of all, it's very ambitious and interesting work. I have just two questions, one on the cyclicality. The

thing I find most surprising is actually the 78 percent versus 29 percent return in terms of net profit --

MS. BARWICK: Right.

AUDIENCE: -- in downturns and upturns. The reason I find that so surprising -- I'm not sure I completely believe it, but maybe you can convince me -- is that -- it's precisely for the reason that you mentioned, which is that if you subsidize people in the boom, you have got a bunch of inefficient firms coming in, whether that's an output subsidy or an entry subsidy -- I mean, they both work the same way -- but then you said that in the downturns they exit.

Well, if they exit and this is a long-run assessment, I find it extremely strange that that doesn't get rectified through the weeding-out process. So the magnitude here I just find implausibly large.

MS. BARWICK: Yeah.

AUDIENCE: Okay, so that's one question.

The other one -- unless the weeding's not working, that's a different story, the weeding out.

The other thing in relation to that is this is in terms of net profit, right?

MS. BARWICK: Right.

AUDIENCE: If you did it in terms of welfare, including consumer surplus in the booms, the demand is

higher, that's going to mitigate to some extent --

MS. BARWICK: Right.

AUDIENCE: -- differential.

MS. BARWICK: Right.

AUDIENCE: The other question is, is it -- I don't know institutionally whether this is the case, but a bunch of these firms presumably are SOEs, right, state-owned enterprises?

MS. BARWICK: Right, right.

AUDIENCE: And I'm wondering whether or not their investment decisions were only based on -- only influenced by the price effect of the subsidy as opposed to dictates, implicit or otherwise, from central government, is the other question.

And the last comment -- but this is more me kicking tires of things that have always bothered me about these kind of models -- which is you assume -- you slid over it -- but you assume that the Markovian process driving the perimeters here is the same before and after. For small changes, that might be okay, but I'm worried, if you have such a large -- a large intervention, whether that's a plausible assumption.

AUDIENCE: Okay. So there are many, many questions. Let me try to say a few remarks.

First, the magnitude. What I didn't mention is

a large number of firms are -- were idle during the recession, and if the bad firms don't exit, they're basically paying the facility without much production, while the good firms are actually producing. That makes a huge difference. I didn't really get into that, about the -- just the sheer cost of being idle, and more than 50 percent of the firms were idle in the downturn. So that's a big, important consideration.

Second, in the ex -- during the booming period, many of them actually were capacity-constrained, and you have this very, you know, quite steep marginal cost and also investment cost, but none of that were true in recession, okay? That's the second very important part.

And thirdly, as I also argue, the selection of different firms, and some of the -- the bad firms, they are not yet exiting. They are just idling, not taking orders, but they are still paying the cost of being active, and the good ones are the ones that can benefit from -- now, the good ones will basically get a bigger part of the market and then benefit from the scale of economy.

So we were struck by the difference, but this is actually an area where we are doing extensive robustness analysis. I believe in the numbers. I can

even give you the, you know, breakdowns of those.

Second, about -- what's the -- oh, okay. What's the last question you asked?

AUDIENCE: Are there possible other channels through which SOEs are investing as (off mic) by the government (off mic)?

MS. BARWICK: Okay, sorry. So there are -- let me go to Markov first. So I completely agree with you that we -- that's a very strong assumption. What we did is allowing the transition process to be different before and after the crisis are the simplest things to do.

Second, we also use different discount factors which -- to proxy for the fact that maybe this is not a permanent -- so if you use a very high discount rate, then firms are essentially acting more myopically, and the quality of answers we get are very similar, okay? And, therefore, the investment decision -- indeed, you might worry that SOE firms are not profit-maximizing. So here we did robustness analysis.

Remember, I said the expansion happened across the board, not just SOEs. So when we look at private firms and joint venture firms, we still see this massive investment expansion, which indicates that this is, again, happening -- you know, basically boost by

the policy of government.

And then we did -- remember, also estimating basically just to say let's now limit our sample, you know, to the private and joint venture firms where we have a bigger confidence in their profit maximization assumptions, and we find similar findings, okay?

The other questions, I will explain...

(Applause.)

MR. VIOLETTE: Okay, great. We are going to take a short break now and be back at 1:50, which is a small update to the schedule. Thanks.

(A brief recess was taken.)

MS. DUTTA: Hello, everybody. My name is Antara Dutta, and I'm an economist here with the FTC's Bureau of Economics, and along with my colleague, Ted Rosenbaum, one of the co-organizers of the conference. So, thank you all for being here.

So after those two great keynote speeches, we are going to return now to the paper session format, and as with our paper session from this morning, we will have two papers presented in this session as well, followed by a discussion of each paper and questions.

So the first paper is going to be presented by Umit Gurun of the University of Texas at Dallas, and the paper is "Non-compete agreements in the financial

advisory industry."

MR. GURUN: Thanks so much for having me, and thank you for having the program organized this nicely. It's so nice to see so many scholars, you know, just interested in the same topic. So this paper is with Noah Stoffman and Scott Yonker from Cornell and Indiana University, and my name is Umit Gurun.

So we started this project, you know, just from this idea that, you know, just like human capital is one of the most important inputs to our production function, you know, starting from Becker, we know that, you know, it's critical for allocating the skilled labor to, you know, most productive and efficient use.

We also know that, you know, just -- it is related to growth of the industry, it's related to the growth of the region, so there's plenty of work to suggest that, you know, when you let people move around, you know, just they will find the best possible way to sort of unleash their creativity.

So having said this, you know, we also know that there are legal constraints for human just mobility, you know, one of them being the noncompete agreements. You know, noncompete agreements are the legal constraints, basically just, you know, these are important impediments to such reallocation, and, you

know, this sort of prevents people to -- you know, to just get stuck in a place that they may not be as useful or productive, okay? So that's one (indiscernible) of noncompete agreements.

The other view says, look, you know, we have noncompete agreements because by doing so firms can invest in their employees and, as a result, they will be more likely to sort of like, you know, just put more thought about, you know, who to hire, who -- to sort of develop the human skill-set, and the result, you know, like that would be a better outcome. So this debate has been around for some time, and, you know, just -- and there are -- there is, like, sharp increase in the use of noncompete agreements.

So just look at certain surveys that we have seen recently, 18 percent of the employees report that, you know, they are somehow bound by the noncompete agreements, and 38 percent of them report that they have signed one of these noncompete agreements in the past in some capacity.

So this is not really something that applies to, like, you know, high-tech workers or the doctors or, you know, the information technology or the high skill-set people, you know, you see this like, you know, with the barbers, you know, the way that, you

know, apparently they cut the haircuts, you know, just probably have something proprietary.

You know, it happens with sandwich-makers, you know, this just made the headlines last year. So somehow people would -- the way that you make the sandwich at certain places, Jimmy Johns in this case, you know, it's proprietary. So if you learned how to make a sandwich at Jimmy Johns, you cannot go to another place and, you know, prepare, you know, another sandwich. So you have to wait a couple of years to forget about that, basically.

So that's basically -- you know, just like it's prominent, and it's like going up a lot, you know, just somehow there is this, you know, just notion that noncompete agreements are necessary for, you know, companies to sort of, like, just keep their know-how somehow, but this increase is becoming important.

And the problem that we have noticed is that, you know, the empirical work so far has been focused on the variation -- in the variation in the state-level enforcement of noncompete agreements, meaning, you know, that we want to sort of identify what's going on, like what's the effect of noncompete agreements in a certain aspect of business.

So you need to find two states, one of them's

enforcing it, the other one not enforcing, and so the one that's enforcing is going to show a certain pattern which is not observed in the other one, so that's like -- you know, that gives the -- that's the core of the analysis most of the time in the noncompete agreement papers.

So what we have sort of unique in our paper is we have the noncompete agreement, the enforcement variation within a firm, meaning after a point, firm says, like, I am going to enforce it, you know? If this happen to leave my company, I will go after you.

And then, you know, there will be some sort of legal consequence of it, and after a certain point the firm says, look, you know, at this point, you can leave, anybody can come to me, so, like, in a way, the firms relinquish their right to sort of go after someone who leaves the company without any kind of -- just like prior warning.

So this, like, within-firm variation is a unique experiment, you know, and this is in a very large industry, say an industry of 500,000 people, financial advisory industry, so these are the small shops that you see in shopping malls, you know, just Charles Schwab, like all these places who most of you may be using, or at least there are 500,000 people

employed in these companies or in this industry, you know, they sort of -- once you get into one of these shops, you sort of are stuck there in the sense that you may have a big list of clients, but you cannot really leave the company and not get sued, because the idea is that when you leave from Company A to B, you will carry your advisors -- your clients with you, you will carry all those assets under management with you, and that will sort of like, you know, be the reason why you will be sued by your original company. So that is the reason. You know, without your clients, you are nobody, okay? So the debate is somewhere in between, okay?

So when you go to a company, you work for, like, Charles Schwab, for example, like you -- the company claims, look, you know, the reason why these clients came to you in the first place is you use my name, okay, I advertise for you, so once they come in, you know, just they are there because of my name, and -- but you're the -- the counter argument is the employees say, but, yes, they came for another reason, but I've worked with these guys for ten years. I sort of helped them to update their portfolio and so on, and so -- so in the very -- like after the tenth year, it's not clear whose customer is this, okay?

When you move from A to B, you know, just some of them may come with you because they think that you, as a person, provide the service, but the -- you know, just -- I mean, it's basically 50/50. So, like, it's not a very clear debate. So (indiscernible), what happened was, after this event, which we call the protocol, in 2004, there was a very harsh, like, noncompete agreement enforcement in this industry. If you move from A to B, there is definitely going to be a lawsuit the day after you leave because they want to make sure that you don't call your clients, okay, to sort of, like, force them to go from A to B.

So the clients still can stay in the original company, but when you move, at some point you can approach your clients and say, look, I'm leaving, you know me, you like me, like I will take care of you, like why don't you sign these forms, and in six months, like, I will take over your assets and, you know, just I will allocate it based on, like, the new company's products.

So given this, you know, prior to 2004, there were like -- you know, lawyers were very happy, okay? Whenever somebody moved, there was a litigation. And so in 2004, some of the companies get together and said, look, all we are doing is we are making the

lawyers happy and rich, okay, so why don't we just sign an agreement, a protocol, which started with, like, four biggest companies, and if anybody sort of moves within this protocol or firms which sign this paper says, look, you can go to another company, and while you are leaving, you can also take the client list with you, and you can call them, and if they want, okay -- obviously, like, you and I both approach at same time, and if the client leaves, fine. If they don't, that's also fine, but you can legally take the book with you. You can carry your clients with you, which also means that you can carry the asset under management with you, okay?

So this makes it really interesting. There's a choice for the firms to sign this, obviously. You know, they make the pros and cons calculation. They say, look, you know, just some -- some people leave, some people come. On average, I will be better off. So, like, you know, just -- there is another consideration here that will be important for us, so -- and this is the event that we are using, okay?

Now, it's very easy to get into this protocol. So you just sign the paper, and in a couple of days, you know, just you will be part of it, and after that, everything is fair game. And you can also sign out of

it, okay? So, again, this is a -- at your disposal. You can say, look, it's not working for us, so we are just going to leave that protocol, and after that, you know, you will be back to old regime.

So this within-firm variation is really interesting for us, and over time, these are the incrementals. So it started with four firms, and in 2005, (indiscernible) one additional firm got in the game. In 2006, another ten, another 18, another 71. As you can see, like, you know, just, like, over time, bigger and bigger, you know, just share of the industry became part of this protocol.

So there were firms that were out of it. There are firms that are in the protocol. So this became one of the major forces in the industry to sort of -- that determine the turnover, okay?

So this is our starting point. So what we would do with this, you know, just -- we study basically a couple of things, okay? First I am going to show you, like, something that labor economists -- I'm a finance person, so, like, first we are going to show this did sort of change the turnover, meaning people moved from A to B. So it's not like a nonevent, okay? So I will show you some evidence on that.

And number two, I am going to show you things

that finance people care about, which is how do money move around the companies, okay? So, like, do we see asset flows going into certain files just because investors are return-chasing? So that's kind of, like, one income criteria in the finance, so just we look at the returns, and that determines our self-worth, or are there other things, such as this one, that, you know, (indiscernible) are totally inactive, they do nothing all day, just like follow the advisors, and if the advisors move will determine that asset flows in the -- sort of in the (indiscernible), okay? So that's kind of like one of the other variable set we are going to look at.

And we will also look at, like, what happens to the fees, okay? So overall is if this is something costly, one thing that we don't observe is the salaries. Obviously if I move from A to B, you know, the salary should be high enough or the wages should be high enough for me to go to the other part. So we don't observe the salaries, but we think that, you know, just it must be the case that this creates some sort of a cost for the companies.

As a result, so our sort of uncertainty is, you know, it should be either internalized by the company or it should be passed on to the consumers. So we are

going to study the overall fee structure by the company to see whether or not they increased their fees so that the customers pay the ultimate price of this, you know, just new environment, the elimination of noncompete agreements, okay? So that's our game plan.

So as I told you, for those of you who are, like, kind of like empiricists, you know, just you are always like, you know, just looking at me, like, isn't this endogenous? Yes, it is, okay? So if you do this, you know, this analysis, basically we have a variable called turnover, meaning you are going A to B, okay?

So this flips to one if your company is in a protocol. That also means, you know, just -- you know, you were in a regime where the company was not really going -- was after you, like, you know, just there was a huge barrier to move. Once this goes to, you know, just -- once you get into protocol, this cost goes down to zero. So we expect to see some sort of action on that beta 2 parameter, okay?

Now, we think this is a little bit exogenous for the firms which are large, meaning if you are working for a big company, such as Charles -- by the way, you know, Charles Schwab is just one name I am just randomly giving, okay? It's not like they are a good example or bad example. It's just an example,

okay? So if there's someone here from Charles Schwab, don't come after me, okay?

So if you happen to be in one of these large firms and if you're an employee in one of these companies, this is an exogenous thing to you because you didn't choose. Your company somehow chose to be part of this protocol. As a result, what will happen to your decision, okay? So your firm chooses something, and if you're one of these employees who is working in one of these large companies, we will assume that this is exogenous and that this is going to give us, you know, a parameter that we can work with, okay? So that's our assumption.

So -- and then, you know, just -- we are going to sort of ground this experiment, but we are also going to look at this in two types of firms, okay? One of them is are you in a state that has noncompete agreements that are enforced, okay? Are you in a state in which they are not enforced, okay? So this is also going to give us, you know, an idea about whether or not, you know, just company chooses this, but at the same time, if you are in a state where this is not enforced, you know, it doesn't really mean anything. So this will be another dialogue that we are going to work with, okay?

So, I mean, the resolution may be pat, but I will tell you and then you will trust me about this, okay? So what we are finding here is that there is no increase in total turnover, okay? What you see is the turnover increases within the firms that sign in to the protocol, okay? So which also means that as firms become more strategic, as soon as, you know, just you are in the protocol, you are going to move, but if you are in a nonprotocol firm, there will be no change in the total move of company, okay?

So that also suggests that this event was, you know, was taken -- sort of is really lowering the cost of the move. As a result, the number of advisors sort of going from A to B in response to this signature of the noncompete agreement, their elimination.

So -- and then based on our expectation, we expect this result to be more strong for the states that enforce the noncompete agreements. We find that that's the case, okay? So that's the -- basically the first column compared to the second column, so which also, again, means that, you know, protocol matters less in states that do not enforce the noncompete agreements.

In other words, if you happen to be in a state that noncompete agreements were not useful -- you know,

were not enforced, you know, taking this out won't really change the environment in that company. That's basically the main finding.

So -- and if you sort of like, you know, try to sort of see this graphically, the year that you sign the protocol is a time that you see the highest turnover spike, okay? So just -- people start moving in and out of this company, just like -- you know, and following the first year, you know, an amicable agreement is reached, so you are back to status quo.

Now, I should sort of like to give you some other sort of background story here. There is a link between this finding and also creation of the small registered investment advisors. So what it means is this. You have been working for a company for so long, you do have all this book with you, and the moment that your company signs this agreement, you tell your brother to, you know, create a small company on the side, and then to have that company sign the protocol, and after that, you move from your company to the other side. So it's kind of a cash-out. And after that, you sell your book to someone else or you retire, okay?

So this also sort of, like, relates to this creation of new companies, so some sort of, like, relationship between elimination of noncompete

agreements can lead to creation of new company. So that's a finding that we have in the paper at the back-end, but overall, the message that I want you to get is that the turnover sort of in the protocol increases, but the turnover in the -- for the companies that are not in the protocol seems not to respond to this event.

All right. So there's also one other sort of, like, nice experiment that happened, like, right about the time that we were about -- you know, just to release the paper. Morgan Stanley, one of the, you know, original signatures -- so firms said, look, you know, if you don't like what's happening here, people seem to be leaving us, we can't keep these people with us, so it seems like I am the loser here. So I am just going to withdraw my signature, so, like, I'll be out of the protocol.

So as soon as they said this, okay, so those spikes that you see -- and the other company also, UBS, okay, they said, like, this is not really working for us. People are leaving. I am the net loser. So ten days before -- remember, there was a ten-day period, like, after you sign, the protocol would be in effect in ten days, people start leaving the company, because you won't be able to leave, you know, just after that

period. So this is like, you know, the closest that you can get in terms of event study with respect to, you know, these types of announcements.

So the company announces that it will be out, and then people start leaving the company, and there is no way to attribute those spikes to some other event happening around that time.

So we think, you know, the evidence so far is consistent with the interpretation that, you know, noncompete agreements are binding, and when they are binding, when you release them, it will affect people's moves, okay, so that's basically the first finding.

Number two, so I'm getting back to this idea of asset under management, as I said, like for those of you who are related to the finance/mutual fund literature, you know, it's a very big question to understand, you know, like what makes people invest and where you invest, okay? So the commentary is you look at deposit returns, and you think that whoever did best, like well in the past, is going to keep, like, doing better, so deposit returns is a big contributor.

So we want to understand whether or not, you know, just the -- the movement of the people really affects the size of the company, because with financial companies, the assets under management is the size of

the company. You are basically what you manage, okay? So you can think of it as kind of a regression of size, regressed on the change in employees that you are getting as a result of protocol, or you can also read it as a -- the asset flow, okay, change in size of asset flow, okay?

So our finding is essentially saying, you know, just -- for one person increase in the -- in the new employees that you are getting after you sign the protocol, you get about -- I have to -- about, like, 16, 17 basis point increase in your total asset size. So, like, your total asset is a hundred billion dollars, here we get about, you know, like 1 percent, so 16 percent of, like, $1 million just attributed to these people coming to you. So if you happen to have 10 percent increase, that will be -- you know, you will get about 1.6 percent increase just because you hired these new people.

So these customers or these flows that are coming to you are not really coming to your own performance. They are coming to you because you are able to attract, you know, just advisors from different companies. So that's basically one finding.

And then -- so there is another finding about the misconduct. So there is a database held by the

FINRA, so we can identify how much of a misconduct is going on in the financial advisory world, so we can see, like, whether there's an advisor who seems to be, like, engaging in some sort of, like, wrongdoing with the customer. We have a database of these complaints.

So what we see -- you know, what's the results we have here suggests that when you become part of a protocol, it becomes really hard to punish the bad behavior. So in the past, when you have the noncompete agreement in place, it would be at a high cost, right? So in a way the company has the negotiation power. When the company has the negotiation power, it can penalize, it can punish the bad behaving advisors, but in the other environment, our interpretation of this result is that the amount -- you know, just people can live at their will.

So if you think that you have negotiation power, but once you let it go, the -- if they see some sort of misconduct, you are more likely to sort of keep that company -- person in the company rather than fire them away. Okay, so this is one of the bad sides of -- a kind of a negative consequence of, like, eliminating the noncompete agreements that we are documenting here.

So the last finding that I want to sort of talk about is this -- the fee. In the asset management

industry, 1 percent is kind of the norm. So if you have -- and we fear one of those, like, people who haven't, like, been investing in indexes, you are probably paying 100 basis points, which is huge, okay? So if you are already doing that -- okay, so one thing that you should do is go and get a index, because it's like seven basis points.

So -- but if you happen to pay 100 basis points and if you happen to be a member of these companies, once they sign in to the protocol, it seems like they increase this 100 basis to 115 basis points. So there is a 15 percentage increase over that, like, 100 basis points, which we think is kind of interesting. Assuming that no other cost structure-related event is happening around the same time, you are basically passing on these costs to consumers or the investors at large, okay? So that's basically our reading of this event.

So remember, like, we are not trying to make any kind of (indiscernible) assessment here, because we do not know how much additional wages these people are gaining. That data is not available to us, but we can say something about, like, the customers are worse off from the fee structure.

To the extent that you think the misconduct

results we have, meaning tolerating the misconduct, it's kind of a, like, negative outcome, this also suggests that as an investor, you are more likely to face with an advisor who seems -- you know, who probably has some sort of misconduct in the past, but he was not penalized in some form. So those are the negatives for the consumers, sort of like lowering the moves in the industry.

So having said that, okay, so, like, again, you know, just this is sort of a -- like, this is a trust industry, okay? So if you want to generalize, it's 500,000 people, like the whole thing between the client and the company is trust. So if the trust is to the company, that's one thing, but if the trust is between the advisor and the customer, then we think, you know, this event is really important, and, you know, it has some big economic consequences that, you know, were not studied before.

So we think, overall, you know -- just, like, I don't want to repeat myself here -- overall, as I mentioned, you know, just -- if there's an economic event that affects about 20 percent of people, up to 40 percent, in terms of, like, you know, signing a contract which binds you to -- not to in the future, you know, we think that's an important event that we

should pay attention to.

So, thanks so much for having me.

(Applause.)

MS. DUTTA: So, thank you, Umit.

So I'd like to welcome Matt Johnson of Duke to discuss the paper.

MR. JOHNSON: Okay, great. Really great to be here. So I'm Matt Johnson. I'm at the Sanford School of Public Policy at Duke. I'm mostly a labor economist, so I'm a little bit outside my element here, but hopefully that means I can provide, you know, some interesting alternative -- you know, adjacent perspective on this really cool paper by Umit and his co-authors.

So I'm also glad that Umit put up that picture of the New York Times article from, I think, 2014, that was about how noncompete agreements were being -- were being increasingly used in hair salons. That article was actually the article that got me interested in this question. I was in grad school. I read that article about barbers signing noncompetes, and I had just taken my grad labor sequence, so I thought, this seems weird. So reading that kind of spurred me to be thinking about this, and it's led me here. So that's been really great.

It actually turns out, a lot of the things that make noncompetes interesting in the hair salon industry is also kind of related to this setting of the financial advisors. Okay, so just this setting that Umit told us about, all right, so here I'm just going to sort of lay this out here, which he already said, all right?

In August 2004 -- oh, I think I'm -- I'm really sorry -- three of the largest brokerage firms create this protocol for broker recruiting, right? So this is in some ways kind of a voluntary relinquishment of a noncompete agreement. So if I choose to join the protocol, I am relinquishing my right to have my workers sign a noncompete.

So, you know, I'd argue this is a really compelling setting to study the effects of noncompete agreements on markets, both for consumers and for workers. This is still kind of a topic that we're learning about theoretically and empirically.

So why is this compelling? Well, you know, some of the main inputs in the production in this industry are what we might call transferable assets, right, assets that the worker could take with him or her upon leaving, the most important of which -- which Umit talked about, right? -- is the client list. And

so if the book of business is something that one party in the firm could take with him or her, that's transferable.

Another one, which didn't get too much -- you know, you might argue that general human capital is an important input into production, right? Investment advisors have to spend a lot of time and money doing -- getting certifications and training that are general in the sense that they are skills that can help me in my current firm but also, if I go elsewhere, I can take those skills with me and be productive there as well.

So what role does noncompetes have here, right? So the -- joining the protocol, you know, by relinquishing the noncompete, the noncompete it -- effectively or relinquishing the noncompete effectively assigns the property rights to these assets to the worker, right? So how do we think about this with noncompetes, right?

So this kind of harkens back to a lot of classical literature. One thing people think about with noncompetes is that they might solve a holdup problem, you know, to the extent that, absent a noncompete, we don't -- we haven't necessarily figured out who has the property rights to the client list.

If I'm the employer, I might think, well, why

would I bother investing to attract clients if I know that my worker can just, like, take them with her, right? So there might be an investment holdup problem if I actually have the worker sign a noncompete. We might solve that by enhancing my incentive to make those investments, making everyone better off.

On the other hand, if it's actually the worker that needs to make the investments to attract clients, the noncompete may hinder the worker's willingness to make those sorts of investments to attract the clients, right? So depending on who's the one making these investments, the noncompete could have different effects.

You know, another thing kind of absent that is, you know, this is something that's especially important in innovative industries, but I think it's probably also relevant here, is that noncompetes are, by construction, going to limit job mobility, right? To the extent we think of workers kind of hopping around to different firms is useful, right, for sort of allocating human capital where it should be.

Also, maybe by promoting innovation and knowledge transfer, you know, is something that noncompetes might inhibit, right? So there's a lot of evidence that -- you know, some speculate that

noncompetes are what sort of eroded the market power of the Boston area's high-tech industry and had to move to California where noncompetes are unenforceable.

So what's the reason why noncompetes are really compelling for this setting? Okay, so just an overview of what this paper finds, right? This paper found, like, four of its high-level findings are when forms join this protocol relinquishing their noncompetes, we saw increased turnover, but particularly turnover with other firms in the protocol.

We saw an increase in assets under management, so kind of increases in firm growth, which is really interesting. They found increases in advisory misconduct rates and also increases in consumer prices, namely, commission fees here. So I just thought it would be useful to put these findings in context with what others in this literature have done.

So this is a very small but very much growing literature, so I thought it would be useful just to tell you kind of what others have done in this space. So -- and this helps kind of articulate, too, the paper's contributions.

So one is this literature has been -- one of the main constraints has been figuring out what's our identifying variation to identify the effects of the

noncompetes. So several papers, including some I've worked on and some other co-authors have used variation at state level, noncompete enforceability, as Umit mentioned, so noncompete enforceability varies across states and also within states over time.

So a lot of papers have tried to use that to figure out how do noncompetes affect these various outcomes. There have been a few others. So some other papers use actual just firm use of noncompetes. This is nice because we actually get something that is at the firm level, which is more fine-grained, of course, and whether I use a noncompete or not is endogenous, so it's kind of hard to use that.

There's one really interesting paper by Matthew Gibson that uses a DOJ ruling that affected the use of no-poach agreements in Silicon Valley, but still, other than that, we haven't really -- you know, that's kind of what we've had to rely on in the literature. And then in terms of outcomes, it's mostly been labor folks who are thinking about the effects of a noncompete, so so far we've kind of been looking at the things that labor people care about, right?

So we have a lot of evidence on the effects of noncompetes on mobility, pretty robust evidence that, you know, the enforceability of noncompetes and the use

of them decreases worker mobility. We had some really interesting work on innovation, so the idea if noncompetes restrict knowledge flows and mobility across firms, suggesting they might decrease innovation, increasingly some evidence on noncompetes and wages, so including some work I've done and what some others have done.

So what this paper does, right, so first, this kind of variation we have is really cool, right? This unique voluntary/involuntary firm-level noncompete use and really noncompete enforceability is a really compelling setting to try to estimate what's going on with them, and especially because there's evidence that noncompetes, they might -- you know, they might cost workers, they might have some social costs, but they really do seem to benefit firms, right? If I, as a firm, get to restrict my workers' flow, I get the benefit from that. So the question of why firms would give this up is a really interesting and compelling one. In terms of new questions, this paper can really send us into new things, such as, you know, how noncompetes affect consumers in the markets that firms operate in.

Okay. So one thing, you know, in terms of the identifying variation, right, the, of course, first

thing that comes up is, is the firm decision to join the protocol endogenous, all right? So Umit, I think, acknowledges this, right? Of course, the year that firms join this protocol is one in which they are growing, right, which we see by the graph that Umit showed us, and so their identification strategy gets rid of two sources of bias we might worry about, right?

They get rid of the time invariant bias by using firm fixed effects. They also have -- are able to account for local shocks, but, of course, there could be time-varying firm shocks that are correlated with whether or not I join the protocol, right? So they have this nice way of accounting for that endogeneity by showing that there's variation in firms' use of -- you know, the effects on this protocol by a state's noncompete enforceability, which is nice.

But one thing I just want to, you know, suggest is, you know, rather that -- you know, I'm an applied micro person. I love trying to find exogenous variation, but maybe this is a place where using the endogeneity can actually be to our advantage, right? It seems to me a really compelling and interesting question, why would firms or under what conditions will a firm give up their right to a noncompete agreement, right?

In some ways joining the protocol reveals a cost of the noncompete to firms, right? So one benefit, of course, is if I do have the noncompete, I get to keep my client list. They are my property right. But a cost is if I want to hire more, if I'm growing, I'm constrained in my ability to do so. So I wonder if we can use -- I don't know, you know, like some kind of revealed preference approach, how big do the benefit -- you know, how big are these costs? How big do the costs have to be to get me to voluntarily just give up my benefits of a noncompete?

You know, why -- we see that only, like, I think 4 percent of firms join the protocol. So, like, why are so few firms joining? Can we use the actual selection model to quantify how much noncompetes inhibit firm growth? And I'll skip over this part for the sake of time.

Another thing I just want to say is, you know, one of the things that is especially provocative at this conference is the paper's last finding about the effects of this protocol on commissions, right, on firms' prices, and I think this is really compelling. I also think the paper could really benefit by pushing this and really thinking about how noncompetes affect market structure.

You know, Umit in the paper kind of defined this thing that joining the protocol led to higher commissions, which they attribute to maybe workers having more bargaining power, but you could also imagine that there's something going on in the background of how noncompetes affect market structure, right?

So there's some really nice work by my co-author, Kurt Lavetti, with Naomi Hausman, where they look at how variation in noncompete enforceability affects market structure in the market for physicians, and they specifically look at how noncompete enforcement affects market concentration, and you can imagine, right, if noncompetes were available to us as firms, the efficient firm size might go up.

For example, I'm less worried about referring clients to my coworkers if they have a noncompete, because if there's no noncompete, I might be worried that if I refer my clients to my coworkers, the coworkers will take the clients with them. So you could imagine in some cases -- and I think it occurred and Naomi found this -- that noncompetes increase firm size and firm concentration, and subsequently find that noncompete enforceability, on net, increases physician prices.

So, you know, I think, thinking about how that might be playing out here, would be really interesting, and I think there's a testable question through which they could do this, just asking if this protocol affected things like firm concentration, it affected mergers and acquisitions, did it affect spinoffs and entrepreneurship, and things like that. I think that would be another way to tease out the effects on prices through different -- a different mechanism than was proposed by the paper.

So I think I am out of time here. So I will just skip to the end. So I think it would be -- in other directions, I think it would be really interesting to look at spillover effects, including if my neighboring -- you know, if one of my competitors in my same labor market has signed the protocol, does that affect me?

There's a lot of evidence that noncompetes have spillover effects by reducing labor market churn and other things. So this would be interesting in its own right. It also might be a way to get around this endogeneity of when firms sign.

But overall, I really enjoyed this paper. I think it opens up a lot of new questions on a very important topic. So, thanks.

(Applause.)

MS. DUTTA: All right. Thank you for that discussion. I am going to welcome Umit back on stage to take a few questions.

MR. GURUN: Thanks, Matt.

Yes, please.

AUDIENCE: It actually was on Matt's slide, but I was looking at your Table 4 when you sort of say state enforced NCAs versus not, and it's binary, and so Matt has used sort of -- I mean, it's a mix of statutes. So is your -- is your 01 just including, like, California and North Dakota and maybe Oklahoma, or do -- because you have sort of the (indiscernible) index, which also relies on judicial treatment, red letter, blue letter, blue pen, or whatever.

MR. GURUN: So we did it two ways. So I think it was in one of the appendices, so we interacted with one of those indices to see as it goes up. So we did it that way. So we thought this was regression -- kind of putting side by side, like present it like that, but we can also, like, bring it up. That's a good question, so...

Yes, ma'am?

AUDIENCE: Originally the broker protocol was among a small number of firms, horizontal competitors,

and they declared a truce among themselves about enforcing noncompetes and defining the terms of -- under which brokers could leave without prompting a lawsuit by the firms.

I don't think that necessarily assigned all the property rights to the brokers, because the firm still had defined the nature of those property rights around them. What the original firms apparently did not anticipate was that other firms would want to join the protocol, and I think probably for antitrust reasons, they couldn't really keep people out of it. So lots and lots of firms joined the protocol.

I think one way of looking at this from an antitrust point of view -- and I wonder if you've thought about this -- is that the original horizontal agreement fell apart in two ways. Number one, more firms came in, and with the number of firms, the intended outcome of the original agreement no longer really served its purposes. And also, there was a falling out among the original horizontal agreement firms as to exactly how to interpret the protocol.

Two phenomena appeared to have happened just by my reading of the press at the time. One was that the original firms started using the protocol aggressively to poach brokers from each other, but to punish those

who left and use the terms as strongly as they could. And the other is that brokers left the large firms and went to smaller firms, and I don't know that you got to that in your analysis.

What happened to the dispersion of brokers across firms and types of firms? How did the relationships between the brokers and their firms change in terms of their compensation and also their other terms of employment at those firms?

MR. GURUN: Can I have my slides back up so I can show something?

So that's a very good question. So you are correct that the original intention was different, so there's definitely an unintended consequence here, which I mentioned in my presentation, that the more and more -- so I mean -- so if you remember the table that I showed you originally, like the size of the firms at the beginning, over time, the smaller firms started to join to the company. So at the beginning, if you look at the number -- like the number of employees, like signatures, it used to be like in the thousands, but later on, so just -- that's fine.

So the -- by the end of -- like, after 2013, '14, the size of the company, like the employee size was, like, in the fifties, you know, like 60, which is

sort of like telling us that -- just, like, it became different. So you cannot always, like, have the big firms sort of be part of this.

So another thing is like this -- so originally it was to approach the people who are in the product of (indiscernible), but later on it was used as an advertising tool to hire new people. So one test that Matt was -- like, when I was, like, thinking about what Matt was saying regarding, like, you know, why is it useful right now, so if you want to be part of this -- this experiment, it has to do with, like, who can you hire, okay?

So we are after the college graduates, new college graduates, so if you are really getting into this industry for the first time, this is really something that you would consider, because this gives you a payout eventually. So when you sign the agreement, it becomes a very nice recruitment tool. So right now, the nature of the agreement is a little bit different. It's used for a different purpose, but originally it was intended for (indiscernible). So that's definitely true.

Now, the one thing -- two other important issues that we don't find, which is whether or not the lawyers -- we definitely know that lawyers are not

happy with this, okay, that's number one. Number two, maybe, like, you know, just they are -- they want to get -- you know, they just can't find the different way of, you know, just like making (indiscernible) for the existing firms, you know, if the costs are not high as much, so that could be another thing that we may be changing over the last 10, 15 years.

And then one final aspect of this is there's actually very active transfer market for the advisors. So you don't really necessarily need to poach me. So if I let people know that in the background, like, if you are producing hundred million dollars, meaning if you have assets under management of hundred million dollars, you may get 1 percent of it if you come to me, but if you stay in other company, it's kind of a transfer market, okay?

So you know that the transfer is giving this much (indiscernible), so that is sort of known to the -- known in the industry, in other words, okay? So there is some sort of like, you know, just competition going on in the prices which we can use, but it's not -- it's not based on the people, you know, just it's based on the company.

Now, on top of that, you may get more or less based on, like, what kind of team you are bringing in,

so...

Yes, sir?

AUDIENCE: (Off mic.)

MR. GURUN: That's why I'm not showing it on the slide, that's correct.

AUDIENCE: (Off mic.)

MR. GURUN: So their argument was that they are not getting as much out of this protocol anymore. So it is possible that this -- we may see a lot of exits from this protocol going forward. So, I mean, it's one of those coalition-buildings, right? You know, just how does the coalition -- how do the coalitions form and, like, how do they break out? So, like, it is possible that in the next three, four years that this variation will -- you know, we will observe the opposite, you know, direction. That's very possible.

So, Matt, thanks so much for the comments. I mean, I sort of like talked about this advertising aspect of it, but all the other stuff is very useful. So market structure, we know that there are new firm creations, so that's important, I think. I think we should sort of highlight this a little bit, but in terms of like -- because this market is very local, so, like, where I live, in Dallas/Plano area, there are a certain number of, like, broker -- you know, brokerage

houses, so brokerage sort of like places. So we haven't seen any kind of, like, big change in the M&A at the local level. So if you can find that, you know, that would be really useful. So there is that local aspect that could be useful for us going forward.

Thanks so much again. I appreciate it. Thank you.

(Applause.)

MS. DUTTA: Thank you, Umit.

So I am going to move on to the second paper in the session, and I am going to invite Matt Weinberg on stage to present his paper called "Oligopolistic price leadership and mergers: The United States beer industry."

MR. WEINBERG: I should have talked to the other speaker. It's kind of the same problem. Okay, here we go.

All right, thanks for having me. Thanks to the organizers and for the people from the Tobin Center for putting this on. This is joint work with my co-author Gloria Sheu, who's sitting right up there in the front row. She may have an opportunity to weigh in during the Q&A. If she does, you should not believe that her views necessarily reflect those of the DOJ or the Federal Reserve Board, her current employer. It's also

with Nate Miller.

What we are trying to do in this project is to provide an empirical model of tacit collusion and use it to quantify the impacts of tacit collusion on firms and consumers, and then to use it as a framework to study prospective merger analysis.

It's motivated by the observation that sometimes you'll see firms adjust their prices in a way that could be kind of hard to reconcile with competition and reasonable changes in costs or demand. An example of this is the U.S. beer industry, spanning the period before and after the Miller-Coors joint venture. What I have in the slide is a figure from an earlier paper that I did with Nate that plots the average price of the three best-selling beers in the United States, Bud Light, Miller Lite, and Coors Light, which we were kind of following relative to inflation for at least seven years. That changed suddenly after the Miller and Coors joint venture. You see that the prices of all three brands jumped up by about 6 or 7 percent and stayed high over the rest of our sample period.

What we did in that paper is we showed that, under the maintained hypothesis of static Bertrand competition with product differentiation over the

entire sample period, given our demand estimates, there were two ways to reconcile that. One was a pretty big increase in the marginal cost of production for ABI brands. In other words, that price increase is bigger than the competitive response that you would get from the standard framework used in merger analysis.

We don't think there's anything going on in this industry that would suggest that ABI had a 10 percent increase in their marginal costs at that point in time, so that suggests a potential -- another potential explanation, which is that the merger changed pricing incentives beyond what you could explain with just the change in a multiproduct firm's pricing incentives given the combination of substitute products.

Okay, so what do you do with that? Suppose that a merger was announced in an industry like that. How do you do a merger simulation, for example? Where do you go with that? That's what we're trying to -- we're trying to provide a framework that you could use to do that with this current paper.

And the way that we approached the problem is just by using the standard theory for studying lawful collusion. There's no reason to think there's explicit collusion going on in this industry, but we're going to

use insights from repeated oligopoly theory to try and provide a framework that you can use to study an industry like this.

There's a challenge in doing that that probably jumps into everybody's heads when I say that, and that's that there are lots of equilibria in these models. There are lots of different ways that you could collude. So it's not clear how to go forward.

In this industry and in some other industries that we describe in the paper, there's a specific pricing practice that we used to justify a particular equilibrium assumption. We call it price leadership, and a good explanation of what's going on in the industry can be found in the subsequent antitrust investigation into ABI's attempted acquisition of Groupo Modelo, the producer of Corona.

What's going on in this industry, as described in that complaint, is that each summer, ABI would make a public announcement of a price increase that Miller/Coors could observe, and then Miller/Coors would typically go along with that. So it seems like ABI is playing kind of a special role in setting prices in this industry according to the complaint, and we used that as a way to justify an assumption on the strategies that firms are using in this repeated game

that we are going to model.

The way it works is that each period ABI, which we view as the industry leader, is going to make a proposed price increase to Miller/Coors. They are going to do that in the interest of maximizing their profits but subject to the constraint that the firms wouldn't want to deviate from that, undercut this suggested price, get temporarily high profits, but then given the strategies that we assume, kill cooperation going forward.

Our framework is going to preserve a lot of the nice features of the basic differentiated product between unilateral effects framework. We're going to allow for asymmetric firms that can differ freely in their cost functions, their marginal cost functions at the product level and their demand functions, and we're also going to allow for partial coalitions within our framework.

We apply it to the U.S. beer industry. We're going to estimate the model, use it to quantify the implications of this super-competitive pricing, and then do something that I find particularly interesting, we're going to use the framework to study the coordinated effects of the ABI-Modelo merger that was reviewed by the DOJ. Ultimately, that merger went

through subject to a divestiture of Modelo's assets to a third party that wasn't producing beer at that point in time. So what we're going to do is we're going to study what would have happened had that divestiture not occurred.

Okay, a preview of the results. We find that tacit collusion is raising prices about 60 cents above competitive levels; however, the industry leader would prefer even higher prices than that. They're not able to get it, though, because it wouldn't be in the interest of the other firm in the coalition, Miller/Coors, to go along with higher price increases, basically because they have more sensitive demand in price and fewer products, so they wouldn't want to go along with it. The cost of that would be too high for them.

We then show that if ABI could control the prices of an important competitor, the Corona brand, Groupo Modelo's brands, that would loosen Miller/Coors' incentive compatibility constraint, and they would be willing to go with higher super markups than we estimate. There's a coordinated effect to the merger.

Okay. So I want to give a high-level overview of the model. It's motivated by two key problems that come up in thinking about repeated oligopoly theory in

tacit collusion. The first is the standard incentive problem, that it needs to be in the interest for firms to go along with these super-competitive prices instead of undercutting them and then killing cooperation going forward.

The second comes from the fact that what we're talking about here, once again, is lawful tacit -- it's tacit collusion, right? There's no express communication going on. And there are lots of different ways that firms could do that, and so we need to think about a way that they might be able to coordinate on one of the many different possible ways this could go forward, right?

So what we view price leadership as is a simple -- a simple way that they might do that. Basically, this is a way for us to justify this particular equilibrium assumption that we're making on the strategies that the firms are employing in this industry.

So the way that this works is there's an exogenous -- we're just taking this as fixed as it stands right now. We're doing some work to think about ways to explore this more in ongoing work, but for now, think about a fixed coalition, ABI and Groupo Modelo, and then within each period, ABI makes this

super-markup announcement, this markup above the competitive price that I'm going to call the super-markup. They announce that, and then after that the coalition members and firms outside of the coalition, Corona and Heineken in this particular example we're going to go through, set their prices, then people buy beer.

We're going to make this as simple as possible. So we're going to assume that there's perfect monitoring. Everybody knows the entire history of play. We're going to assume that everybody knows the current cost and demand in the economy. There's no kind of private information here that's different than some models of price leadership that are in the theory literature. That's all common knowledge in our setup.

And what I'm going to show you to motivate the intuition behind what's going on, we're going to assume that the state of the economy is IAD stochastic and unaffected by current actions, and the empirical implication -- advocation, excuse me, if we are going to instead assume that there's been a perfect foresight about what costs and demand are going to do in the future. We're going to view the outcomes of prices and quantities as the result of a subperfect equilibrium where firms are playing these particular strategies.

So, again, the way this works is that the strategies are going to take this grim trigger form, so very simple strategies, where initially they're each going to go along with this markup -- I've called it M subscript T -- and generally can vary with time. I'm going to show you results where we just fix it to be one number.

So they start off by setting prices equal to the super-markup above the competitive differentiated product Nash-Bertrand prices instead of the static version of the model. The fringe firms set their prices with the understanding of what the coalition is doing. So they're picking prices to maximize their profits given the operations of the coalition. And each firm is going to weigh the present value of going along with this against the present value of deviating, maximizing their profits given what the coalition is doing and the state of the economy, and then killing cooperation going forward and getting back to the Nash-Bertrand prices. So that's the structure of the strategies. It has a simple grim trigger form.

So as is standard in the study of repeated games, we're going to think about the incentives of whether or not you would want to go along with this in terms of this weighing of the present value of keeping

coordination going against windfall/temporary profits and then competitive profits going forward. We're going to summarize these incentives with these functions that we call the slack functions, which just is a way to put a number on the difference between the value of staying in the coalition versus killing it, and it depends upon a few things that we're going to try and compute given our estimates of demand and costs and how firms are setting this super-markup.

The first component, of course, is the price leadership profits. Those depend upon the current state of the economy, and they're also dependent upon this cost of coordination. This is this R function that I'll talk more about momentarily. The expected future net benefit price leadership is that going forward net of the Nash-Bertrand prices, and that needs to be greater than the value of deviating relative to the foregone future price leader profits, so it's that immediate net benefit of deviation.

So if firms are going to be willing to go along with a proposed super-markup, then this needs to be strictly positive -- or, sorry, it needs to be positive for all firms in the coalition. By construction, it would be weakly positive for firms on the fringe.

Okay, so this antitrust risk function, why is

it there, first of all? So in the framework of our model, for any discount factor, some level of coordination would be possible in the absence of this cost of coordination. We've put it in there for now. We think that it might reflect the risk of increased antitrust scrutiny on proposed mergers in the future in an industry where this sort of pricing is going on, as was the case in the Grupo Modelo-InBev transaction.

We make a couple of assumptions on this form. We're going to parameterize it in the empirical work. We assume that it's zero if there's no super-markup and you're at competitive levels and it's increasing in the markup, so in the markup above the competitive price.

Okay. So, again, ABI, when they choose this markup term, they're doing it to maximize their profits subject to the constraint that everyone wants to go along with it instead of -- instead of deviating. So that's the form of the model.

We apply it to the beer industry between 2001 and 2011. Over that time period, there were five of the four firms after the Miller-Coors joint venture in 2008. The biggest firm in the industry was ABI. They had about 35 percent sales -- of sales -- revenue, I should say. The Miller-Coors joint venture created the second biggest firm, which wasn't that much smaller

than ABI, about 28 percent of the market.

Corona and Heineken are -- everyone here, I'm sure, knows what these products are, I feel kind of silly describing them -- but these are higher priced imported brands that, importantly for our framework, were viewed as being outside of this coalition. So in that DOJ investigation, they were described as taking a different pricing strategy. They were not going along with the price increases that ABI was proposing.

We are going to focus on the most popular beers in the market. There are 39 different products, 13 best-selling brands, different package sizes, and our data spans two important mergers. The first is Miller-Coors. The second is this merger that was modified with this remedy that required the divestiture of the Corona brands that's going to be the subject of our counterfactuals.

Okay, so I want to give you a brief sketch of how we're going to operationalize this, how we estimate it, and then I'll move on to the results and the counterfactuals. So the first step in doing this is to try to understand what firms' marginal costs are at their product level given this pricing/quantity data and the structure of our model.

So suppose that you know what demand is, that's

estimated. You know who's in the coalition and who's not. And suppose you've got a potential super-markup. What we show in the paper is that even if it's a partial coalition, you can infer what marginal costs are given that information, right?

So the -- it's easiest to say if you think back to kind of the standard framework for figuring out what marginal costs are when you have price and quantity data in differentiated product industries. If you know what Nash prices are, you can compute the marginal revenue for each firm, and they are going to be setting that equal to marginal cost, right?

Here, we don't know Nash prices. We have to get back to them, but we've assumed that the observed prices are the sum of the super-markup and the Nash price, right? So if you just -- in the scenario where all firms are in the coalition, if you just subtract off this candidate super-markup, you get back to the Nash price, and then you could compute what marginal costs are by constructing the marginal revenue from a -- from the static Nash-Bertrand model.

Okay, so different super-markups imply different marginal costs. What we do next is we parameterize the marginal cost function to depend upon some things that we can measure, things like shipping

distances between retail locations and breweries. We allow them to vary freely by j, which indicates the product and across regions, indicated by r, and over time.

What's left over is an unobserved marginal cost component. Marginal revenue depends upon the super-markup and also upon Nash prices, which in turn depend upon unobserved marginal cost shocks. So this markup term that we're going to try and estimate is a -- it's a choice. It's an endogenous thing, and in order to do this, we're going to make the assumption that ABI's marginal costs aren't increasing relative to the fringe before and after the Miller-Coors joint venture.

This is a part of the paper that we're currently working on. I'd be happy to talk to some of you guys about this, about ways that we're thinking about extending this later on, but for now that restriction gives us an equation that allows us to pin down a particular markup term.

Okay, here are some results. They vary across different demand specifications, which are indicated by these different columns. Those are taken from my previous paper with Nate. Columns 1 and 3 are both estimated at the monthly -- on monthly data and differ

in the exact specification of heterogeneity in our -- in our BOP demand model that we used to get, you know, elasticities that we think make sense in this industry.

Columns 2 and 4 are the same thing but estimated on quarterly data. Across specifications for getting this markup above competitive prices to be between 60 cents and 74 cents. We estimate that Miller-Coors -- Miller and Coors' marginal costs fell after the joint venture by about 53 cents and 83 cents respectively, in column 2, for example, and we also estimate that their marginal cost of production in the beer industry depended upon the shipping distance between the retail locations and breweries of heavy, bulky product. That's what that gamma 3 parameter represents.

With that information, we can calculate what marginal costs are for each product instead of each market, and with that information, we can compute the markup that ABI would like to set if they could ignore incentives for firms to not go along with the price increase, and it's substantially higher than what we've estimated. It's on the order of, like, $2.50 as opposed to 60 cents, suggesting that one of the firms is just indifferent between going along with the super-competitive price and deviating -- or implying

that, not suggesting it.

The important welfare implications, we estimate that firms' profits are about 8.5 percent higher relative to standard Nash-Bertrand profits because of this pricing strategy that they're able to implement, and that for every dollar that firms gain, consumers lose about four.

For the rest of the paper, I'm going to show you results from the second specification. It's on quarterly data. It's a little bit easier to do computations where we need to get the rest of the results.

Okay. So one of those firms is just indifferent between going along with the -- with the coalition and deviating. What we're going to do next is we're going to try and learn about the remaining parameters of this model, the two new parameters that we've introduced relative to the standard framework, which are this antitrust risk function, r, and this term delta that you can view as a discount factor.

We try to go agnostic on exactly how to interpret delta. We think it can reflect a number of different things, including, you know, how severe punishments would actually be if there was a defection. Perhaps they're not as extreme as grim trigger. That

would map into lower values of delta if the modeling were specified in that way, or, you know, the chance that the game just ends at some point, and so on.

So one of these has to bind. What we do is we parameterize the antitrust risk function just to be simple, a simple linear thing, and compute the different combinations of delta and the parameter of the risk function that would cause one of these slack functions to be exactly equal to zero, and here the different combinations in this figure, you can see there's a tradeoff.

For higher levels of the discount factor, you need higher antitrust risks, which is a different combination set that would satisfy that equation. Here are those slack functions for the two coalition members, ABI on the left and Miller-Coors on the right. Miller-Coors is the firm that's just indifferent between deviating or not at the estimated super-markup, whereas ABI would prefer strictly higher markups. You just can't get them because Miller-Coors won't go along with them. You can see that if the discount factor were 0.7, they would be able to -- they would prefer a markup of about, you know, just under 80 cents instead of 60 cents.

Okay, so the last part of the paper that I want

to talk about is how you can use this to do counterfactual merger simulation. So what we're going to study is ABI's attempted acquisition of Grupo Modelo. In DOJ's complaint, in their investigation of that merger, it says that Corona, the Grupo Mondelo brands, seem to play a really important role in preserving competition in this market.

In particular, after ABI would announce these price increases and Miller-Coors would go along with them, ABI did not, right? So there are more expensive beers, but the gap in the price between Corona brands and the best-selling beers in the market was shrinking over time because of this. So they were very concerned about the elimination of Grupo Modelo.

Okay, so here is just those -- the plot of those slack functions describing the difference in the present value of going along with the cartel and deviating after ABI's able to control the prices of a key substitute, Corona, and you can see that it relaxes both of them. You can get an estimate of what I'm going to call the coordinated effect of the merger by just looking at when the slack function crosses the horizontal intercept here. You can see that after the joint venture, if the discount factor were 0.5, they would be able to get a super-markup of 73 cents more,

so $1.33 in total, right?

So we can decompose the total price change into two different components, one that I'm going to call the coordinated effect, which is illustrated in that graph; the second of which I am going to call the unilateral effect, which comes through in the model because observed prices we've assumed are the -- the sum of this markup on top of competitive prices that the coalition is setting and Nash prices. So that's the basis for the observed prices, and those are going to change, too, with this joint venture.

This table shows results. You can see that the change in Bertrand prices, the basis -- you know, the difference between observed prices and the markup, like the super -- what we call the super-markup, that end parameter that we've estimated, it goes up by quite a bit. It goes up by the most for ABI and Modelo, and between those two, it goes up for the smaller of the two firms involved in the merger, as you would normally expect, inside of that unilateral effects framework, right, because the competitive constraint for ABI was more important for Corona than the converse.

You can see that the super-markup, you know, if delta is equal to 0.7, goes up by $1.01, so the total price change is pretty big for Modelo, because they get

the entire increase in the super-markup, the original 60 cents plus the bump-up, plus the change in Nash price, so it's about $3.36 if delta is equal to 0.7, which has a really big impact on their market share. It goes down by about 50 percent. So this supports the DOJ's theory that Corona was very important for preserving competition in this market.

Okay. The final thing I want to show you is the implications of efficiencies in an industry that works like this. This is kind of the classic thing that you do when you review a merger. You've got -- you know, in the unilateral effects framework, you've got this upward pricing pressure that comes through multiproduct firm pricing incentives, and you want to know, are marginal cost reductions sufficient to outweigh the incentive to raise prices from providing substitutes?

So we consider three different scenarios and contrast what you would get in the stated unilateral effects framework with what you would get in this model of tacit collusion, what we call the price leadership model. The first scenario is one where there is no marginal cost efficiencies; the second is one where Corona experiences a 50 cent marginal cost reduction, we call that a minor efficiency after the merger; and

the third we call a major efficiency, and these are the marginal cost reductions that are exactly what would offset the increase -- the incentive to change prices at all. It's about 51 cents on average for ABI and $1.72 on average for the Grupo Modelo brands.

The top panel here shows the Bertrand prices across the two different frameworks. The numbers differ slightly because in the right-hand panel we assumed that it's a price leadership model. If you read the data on the left hand, it's assumed that the Bertrand model generates the data. The different -- you know, those change kind of as expected. If there's a minor efficiency, the price increase for Modelo brands goes up by $1.15 instead of what you would get if there were no efficiencies, which is $1.70. It's the same in the PLE model. And, of course, if there's this major efficiency by construction, there's no price change in the Bertrand model, and it's pretty close to zero in the PLE model.

What's different, though, is the impact on this coordinated effect, this super-markup, right? So the intuition here is that really what's limiting the ability of the coalition to increase prices is Miller-Coors' incentive to deviate, right, which is not a party to the merger, right? So this -- the

implications of merger review -- of this model from merger review are pretty different than the standard model, right? So the marginal cost reduction is not going to do much to offset this increase inside of the coordinated effect, because it -- you know, what matters is the third firm whose costs aren't changing.

You can see that -- the final thing I want to point out is that, you know, across the different specifications, this merger is worse for consumers under the PLE model than it is under the Bertrand.

Okay. So that was a lot of information, but in summary, what we're trying to do here is to provide a framework for understanding the relationship between market structure and prices in an industry where we don't think that the kind of standard framework can explain a lot of the changes that we've seen in the industry. We've shown the implications of this matter for merger review. The pass-through of efficiencies is really different when there's tacit collusion going on than when you've got a Bertrand competition and that it matters for both consumers and for firms in terms of welfare and profits.

All right, thank you.

(Applause.)

MS. DUTTA: All right. Thank you, Matt.

I would like to welcome Gaurab Aryal, from the University of Virginia, to discuss the paper.

MR. ARYAL: Thank you. Thanks for the invitation to discuss this paper.

I saw this paper and discussed this paper one year ago in D.C. at the SEA, and when I was asked, I was like, huh, it could be easy, but then I think the paper has changed a lot, and so -- you know, in a good way, so I had to kind of re -- try to remember what I was saying then, but it didn't make sense, so I had to redo everything.

So the paper -- I'm just -- I don't have much to say in terms of the content, but -- or what they should do, but I think the paper is fantastic, especially given the fact that they're trying to estimate a structural model -- dynamic structural model with collusion, estimating the welfare and efficiency, and doing everything in a way that makes sense was -- it's fantastic. So I think all of you should read the paper if you want to.

So I'll just try to recap what Matt just talked about and maybe just point to one topic on the antitrust enforcement, okay? So generally -- so the basic summary is that they consider (indiscernible) to be a price-setting game, where there's one leader who

sets Bertrand price plus something. That something is what is known as the super-markup. Of course, we don't observe that, and we think that we would like to estimate that.

And the member firms, in this case pretty well known from their previous paper, if they value the future enough, if the price satisfies the incentive compatibility constraint for them, that's what Matt talked about in terms of the slack function, and implicitly, there is good monitoring, then it provides incentives for them to participate in this price-setting collusion or whatever you want to call it, right? And so they show that there is a sort of imperfection in this relationship.

The application of this is, again, the beer industry. The context is the post-2008 Miller-Coors merger where the ABI is the price leader, okay? So the -- ABI moves their head and signals everybody to follow suit, and they estimate that the markup, which I called "something," is about 6 percent of the price. The price leadership also, of course, naturally increases the profit, decreases consumer surplus by a large amount. So if you were to just add the (indiscernible) surplus and consumer surplus, this seems to be welfare-decreasing, which is not

surprising.

And I'm not going to talk about this because I don't have time, but they also do this in the -- a big chunk of the paper is to think about the ABI-Modelo merger and say something about efficiency and contrast that with the welfare implications, okay? And so what we find is that -- the first thing they find is that constraints are relaxed, which is what we would expect. The larger the coalition, you know, more -- more firms would like to participate, and, therefore, the super-markup increases and amplifies according to effect, right? So I think that's kind of nice.

There are a bunch of empirical challenges. I'm just -- there's a lot to be said about the paper, but I will just focus on a few things. First of all, I mean, as I said, writing a tractable structural model of this repeated game, it's hard. It's really hard, and there are a bunch of subtle modeling decisions that goes through this.

First of all, the coalition itself is not optimal, so they're losing some money, and so as far as I know, from my own work, I don't know how to deal with coalitions that leave money on the table, because -- I mean, it is essentially saying that the first order condition is not really binding, there is some lag, so

it's hard, right? That's one thing.

In terms of the identification, the "something," which is the super-markup, essentially there's one equation, two unknowns, hard to identify, and one of them is, of course -- so one is the marginal cost that we would like to recover and the second one is the super-markup. So what the key assumption is is that before the 2008, the super-markup was zero.

So what that means, that you can look at pre-2008 data, estimate the markup, go to the post-2008 data, then use the markup that you've estimated before, and get the super-markup, because now you have one equation, first order condition, and one unknown, which is the super-markup, and I think -- thinking about it, I think it made a lot of sense to me.

Again, the crucial assumption is that the cost does not change, right? Otherwise, this wouldn't work, and from the previous paper that Matt and Nate have, this seems to be a reasonable assumption. So without that paper, I think one might have to argue why this could be the case. So I think that was -- that was nice.

In terms of -- I just want to also talk a little bit about -- so I was thinking about where it fit and where I have seen something like this, and so

this is -- this is a student at UVA is looking at the retail vertical bargaining for wholesale and retail price, and he observes only the retail price, doesn't observe the wholesale price, and so, therefore, cannot estimate -- has to estimate the bargaining power and the marginal cost, something similar to what they face.

And so what he does is he uses the data from (indiscernible) control state where the retail markup is fixed to back out the cost and then go to the nonalcohol-controlled state to estimate the bargaining parameter, which is very similar in spirit, if not exactly the application, of what they're doing, and there's also this paper by Matt Grennan in 2013 where he is also exploiting something extra. So the point is that if you want to estimate something, you have to have either additional data or different regime. I think that kind of fits pretty well in that thing. So I thought it was nice.

But I want to just point, in the last two or three minutes, I want to focus a little bit on the antitrust risk part of it, which I thought was an innovation. So usually we understand that there is some risk, right? Maybe the risk is proportional to the money that you make from the cartel, but we typically -- at least I have always ignored that risk,

although I'm aware of such a risk, but here I -- it was nice to see it being incorporated in the payoff function in the form of r, which is a function of super-markup.

One can debate whether it's -- you know, the functional form of (indiscernible) separable or multiplicative, where the damage issue -- you know, where the risks or the antitrust risks should be proportional to the damages that are inflicted on the consumers or not, I mean, that's -- I mean, keeping that aside, I thought it was kind of nice.

Because in my -- going back to my class when I'm teaching, so I teach an auction class, an undergraduate class at UVA, and when I'm talking about collusion in auction, I try to make a -- that in contracting there is this risk, but we don't know what to do with it, and so let's look at larger projects where, you know, the benefits at least are good for the firms to collude in. And so I tried to wing it, in other words, so it was kind of nice to see that there is a way at least to have some idea of what this function looks like, right?

And so -- and, in particular, I don't know of any paper that tries to look at the firm's expectation about -- about the antitrust risk and how if affects

the outcome. There are lots of recent papers that look at different regimes, right? And so the crucial point is, how do we identify this r? Again, that's another parameter. This is firm's belief.

So probably a naive exercise could be -- I mean, if -- an ideal exercise probably, even if naive, would be so if we could fix the pie, which is how much the firms can make of this PLE, and if you could vary the antitrust regime, then there is some hope that we might be able to trace the r, but it's -- again, this requires that the antitrust regime changes, which is not the case, and so it's nontrivial.

So here what the authors do is they first "depermatize" the r into a very simple form, which is still not identified, of course, because -- and then they use this idea that at least for the Miller-Coors, the incentive compatibility constraint must bind. As Matt showed, ABI would like to have a higher super-markup that it is able to charge, but Miller-Coors is not going to go, so we have some idea about -- you know, think about it as a moment condition, if you will, and -- but the problem now is that this antitrust now gets mixed with the discount factor, right?

And so now we have two parameters, the phi,

which is the antitrust risk, and the delta, which is the discount factor, both of which are unknown, and depending on -- as we know from the folk theorem, depending on your delta, you are going to have different outcomes, so -- and, you know, identifying delta is a nightmare. We always put 0.95 and go home. The only paper that I know which tries to identify the delta is a paper by Abbring and Daljord, right?

So this is a tradeoff, as you can see, discount factor low, you know, or higher risk coefficient. So I want to just point out what -- it reminded me of -- there's another paper by Perrigne and Vuong where they're trying to identify Lafonte and Tirole's optimal regulation model, where they have this structure, where the regulator is trying to maximize social welfare, which is a function of consumer surplus, plus lambda times the consumer surplus, where lambda is the weight on the producer surplus.

If it's one, then it's consumer surplus/ producer surplus, the same. So what they do is that they -- they try to find the -- ways to identify this lambda, and what they find is that if the payment that the regulator makes is not exactly given by the optimal contract, there is some noise in the process, there is hope to identify this lambda.

So what I was thinking about is that perhaps when you're doing this -- the slack function where you have these states, which is treated as known process -- in fact, in real world, probably they are not known, so there's some noise -- and so possibly there's a way in which you could leverage the variation in the process, the state space or the demand space, to identify or say something meaningful -- more about our -- than you are able to, but, of course, I don't expect you to do that in this paper. So I encourage you to kind of think about this and write a paper for us, right?

Thank you.

(Applause.)

MS. DUTTA: All right. So before we wind down the session, we have the opportunity to have a Q&A on this paper that was just presented.

So, Matt, and I don't know if, Gloria, you want to come up on stage, too. Thanks.

AUDIENCE: Hi. Steven Bristoll, Federal Maritime Commission. By the way, excellent paper. Have you thought of any way to proxy for maybe change in antitrust risk or maybe the perceived change in antitrust risk, like maybe changing -- changes in regulations or maybe if there were, like, a newspaper out -- newspaper articles that said, oh, one of these

regulators is considering a probe into this industry? I would think that maybe then the companies might be -- back off their collusion a bit.

MS. SHEU: So we've thought a little bit about different ways to parameterize that or to add some data to it. I don't -- we didn't -- we haven't had -- we don't have any, like, specific examples that we have actually, like, run through. I mean, it's a good thought.

To a certain extent, this part of the model is -- it comes from the quirk that I think with a lot of profit functions that are kind of smooth, you would -- you know, if you could collude a little bit, you want to do it, right? And part of what we want to generate is the opportunity to not always be colluding. So there's, like, a little bit of a technical part there.

So this can stand in for, like, a lot of different reasons for why you're not colluding all the time. It could be antitrust risk. I think that's a totally natural interpretation for it, and that's one that we -- that we explained in this paper, but it could be other things. Like, it's just logistically difficult to do this, especially if you can't literally call up, you know, or get in a smoke-filled room and

talk to your rivals. So there's -- I am going to acknowledge, there's a bunch of different ways to think about that.

AUDIENCE: So it could be like trembling hand?

MS. SHEU: I mean, it's anything that kind of keeps you from getting to that perfect, like smooth, like I am going to raise my price by a little bit to collude. So there's different interpretations.

MR. WEINBERG: It's a kind of an unfamiliar abject, right? You don't see it in a lot of these models. We have played around with kind of different assumptions on its form, and the results we showed you today, it's just a fixed per-period cost of coordinating that is increasing in the gap between the competitive price and the price that the firm set.

Like some of the stories, you might think that maybe you should also depend upon firm size or how much they sold, like you thought it reflects the risk of damages or something like that, that some of those things should come in as well. It doesn't right now. We've explored some of that right now, and we think that maybe we should do a little bit more of that going forward, because it is kind of a new abject.

AUDIENCE: How much changing it around does (off mic) to affect your results?

MR. WEINBERG: I mean --
AUDIENCE: Or you haven't done it enough or --
MS. SHEU: No, I think overall we still get -- as long as you believe the form of the ICC, that the insider compatibility constraint is roughly what we're talking about, you are going to get, like, the same story. It's -- you know, there's a question of, like, is there separate interest on exactly what's going on with the risk term, and that part we haven't necessarily fully explored.
AUDIENCE: Okay, thank you.
MR. WEINBERG: Thanks.
Yeah?
AUDIENCE: (Off mic).
MR. WEINBERG: So --
AUDIENCE: -- interesting paper --
MR. WEINBERG: Is this on? Okay, so it seems like the existence of the paper kind of undermines the --
AUDIENCE: (Off mic).
(Audio issues.)
MR. WEINBERG: It seems to be --
MS. SHEU: I can hear you, Joe. Just keep going.
AUDIENCE: So -- so it -- there's perfect

information. So if you're saying that in the real world, the antitrust authorities have imperfect information, and the higher prices go, the more likely we are to have our next merger blocked, then we don't want to raise the prices too high, because then they will infer that the market's already very high and there's more opportunity to raise prices, but now with the technology that your paper has given potentially antitrust authorities, now there's no information that we don't know, so now we know that you're setting prices in order to hide the fact that there's market power.

MR. WEINBERG: Thanks for your optimism and your endorsement.

MS. SHEU: I'm sure we're all playing that super game right now. I mean, this is -- again, I mean, this is one way that we think in this -- in this -- in this example it appears to be the way that they're communicating and the way that they're colluding. I think that there's a bigger thing that we all need to think more about and do more research on, which is just, like, generally, what are the different ways that firms can collude?

There's a large menu of them, and I don't know that we necessarily have certainty generically coming

up to an industry that they're all colluding in exactly this fashion. I think we've seen a lot of things that are consistent with the price leadership thing that we have put forth in certain situations. I wouldn't say that that's, like, true across all industries.

AUDIENCE: So your identification stems from the fact that you assume that there are zero super-markups in a preperiod and there was super-markups in a post period. Would this also just be consistent with the super-markups in a preperiod being -- increasing to a new super-markup, or is that an inaccurate reading of the model? Like super-markups before 2008 could have been 60 cents and now they've become 68 cents, or do you require that the fact that there were zero super-markups in the preperiod for the inference to make sense?

MS. SHEU: So I don't -- so I think some -- so I can't be totally precise on this because, you know, there's a lot of moving pieces here. I think that if we're not right that the super-markup is zero in the preperiod, it will affect some of our marginal cost estimates. So I think that it wouldn't be perfectly true that our results would kind of, like, go through exactly, but I do think that, like, some of that -- some of that flavor -- I'd say it's not going to fully

load. I don't think it would fully load on the marginal costs, and I think some of that flavor would go through. I don't know the exact magnitudes, and that's also something that we've been thinking more about in our current identification strategy, kind of like thinking about revisions to the paper.

Our -- you know, that was -- this is our initial strategy that we took, in part due to, I think, like the evidence that came out of Nate and Matt's previous paper, which I think is pretty compelling, but we are, like, exploring other things. I mean, part of the issue -- you know, this came out in the discussion that we just had, which was, you know, you need to -- you have multiple things moving around here that you want to identify.

You have the marginal cost and you have the super-markup, and to -- like, if you -- if you're able to instrument in such a way that you have something else that moves one of them around and not the other, you don't necessarily have to set it to zero. You don't have to set one of those things to zero. So that's an idea that we're playing around with right now to explore this more.

AUDIENCE: Okay, I just had one more question. So have you looked at all to see if the timing that you

assume is consistent -- like, for instance, is there any chance that one of the fringe firms would have an incentive to jump before the price leader and say, hey, you know, we're going to set a low price this period and not be a follower? So if -- essentially endogenizing the timing.

MR. WEINBERG: We haven't done much of that. I think behind the scenes a lot of what's going on is that there's this kind of communication that's going on, public communication, not necessarily in a way that implies antitrust violations, but we think that ABI is kind of making this price adjustment known, perhaps through the distributors, something like this, and then the prices are actually being set kind of simultaneously afterwards.

Everything we've done empirically has been kind of at the monthly level or the quarterly level. We have not been looking at, like, super-granular pricing adjustments over time, but that's an interesting thing to think more about.

MS. SHEU: Yeah. I mean, to a certain extent, the data we're using is at the retail level, which has some additional noise in it for a variety of other reasons. So some of the -- some of the -- the data that we're using aren't necessarily speaking to that

timing at that level of precision, so that's not really something that we've looked at for that reason as well.

MR. WEINBERG: All right, thanks.

MS. SHEU: Thank you.

MR. WEINBERG: Also, thank you, Gaurab, for the discussion, wherever you are. Thanks.

(Applause.)

MS. DUTTA: All right. So we are going to take about a 25-minute break and be back here at 3:50 for the final session of the day. Thanks.

(A brief recess was taken.)

MR. VIOLETTE: Hi, everyone. Welcome to the last session of the day, and we're excited to have Zarek Brot-Goldberg from Yale University presenting his research on intermediation and vertical integration in the market for surgeons.

MR. BROT-GOLDBERG: Okay, thanks for having me at the conference. This is joint work with Mathijs de Vaan, who is at the Business School at Berkeley.

That timer is definitely wrong. It's counting the wrong way.

So the motivation for this paper is in the last couple years we've seen a rapid acquisition of -- so we've seen a lot of consolidation in the U.S. healthcare industry, mostly in -- in recent years,

we've seen rapid acquisition of physician practices by large innovative health systems. You should think of these as large conglomerates that own lots of different parts of the healthcare supply chain, from hospitals to radiology clinics to outpatient facilities and now increasingly physician practices.

So between July 2012 and 2016, the percent of U.S. physicians either in a practice owned by a health system or directly employed by a health system doubled from 14 percent to 29 percent in the practices and from a quarter to almost a half of physicians. That comes with a lot of other forms of vertical contracting which essentially can replicate some of what we might think of as vertical integration.

And it's been encouraged by recent innovations in both public and private reimbursement, including accountable care organizations as well as private COOs, like global risk contracts.

Now, what we're going to think of is that integration allows for coordination between doctors, specifically in this case primary care doctors and specialists, which we'll think of as creating a tradeoff. So on the one hand it may allow for the increase of the productive efficiency of care by allowing for these two parties to coordinate on how

they care for a patient. On the other hand, it may damage allocative efficiency by creating incentives for primary care doctors to steer their patient referrals towards specialists they're integrated with, away from more efficient outside options.

We think of this as roughly akin to standard efficiencies versus foreclosure (indiscernible) we think of in vertical antitrust, although it's a little different because we'll think about productivity rather than prices.

Now, the healthcare context is one where we worry a lot about general productive efficiency and about allocation given sort of dispersion productivity as well as institutions which we might think dampen the ability of patients to sort efficiently cross providers. Along with that, in healthcare as well as sort of I think more generally, despite the fact that we have given three or four Nobels for the theory of vertical integration, the sense in which -- whether we know if vertical integration is anticompetitive in healthcare and beyond is still fairly limited, and in healthcare, we have a number of reduced-form estimates of the effect of vertical integration, but I would say in this paper we're hopefully going to expand that to have a sense of why these estimates have varied so much

across papers.

Okay, so what we'll do today is we'll ask, how does vertical integration shape the productivity of U.S. healthcare spending? And we'll think about what drives the heterogeneity across different sorts of systems and the effect of vertical integration, including these productive efficiencies and allocated distortions, as I mentioned, as well as differences in market structure across firms.

We're going to have a very narrow application, which will be orthopedic joint specialists in Massachusetts. Now, I think this is sort an interesting and representative market for healthcare in certain dimensions but also a little peculiar in other dimensions, so this is an extremely integrated market in the sense that almost every PCP and orthopedist has an integrated -- has some form of integration with at least one of the other party, and two-thirds of orthopedic patients are coming from a primary care doctor who is integrated with their orthopedist.

Now, the fact that this integration is so pervasive, as well as the complexity of these different moving parts of how integration can affect outcomes, is going to make identification of the counterfactuals that we care about very challenging, and while I'm not

going to solve the endemic issues of identification in the vertical integration literature, I will hopefully convince you that the assumptions we need to make in order to get these counterfactuals out are things we may or may not be willing to accept.

All right, so -- all right, let me skip the roadmap. So health insurance -- healthcare is very complicated. I'm not going to have the time to describe all the institutions in detail, but let me at least talk about why we might think that primary care doctors are going to integrate into these large organizations. So the reason that -- there's a number of reasons, but the reasons that folks say is that integration allows for coordination across providers.

Now, why is that true? Well, both anti-kickback laws and health privacy laws restrict the ability of health -- of physicians to coordinate and contract across firm boundaries. So HIPAA, for example, is going to restrict the ability of a doctor to just call up another doctor, even within their organization, to talk about a patient who they may not have records for, whereas being able to be part of the same organization allows you -- allows doctors to coordinate jointly on, say, IT adoption or other forms of communication adoption which will facilitate

transfer of patient records and other forms of communication.

Why do we like this? Well, we like this because it potentially has the potential to reduce duplication of effort. So we worry a lot in healthcare about wasteful spending. The ability to coordinate who does what will help -- we think might make, say, a primary doctor and an orthopedist who I visit don't, say, order the same x-ray for the same patient.

Now, on the other hand, we worry that systems -- these sort of large systems can use incentives to encourage PCPs to steer patient referrals internally. I should be clear that paying for referrals is considered illegal under these anti-kickback laws; however, we've seen from a number of cases that systems are able to use other payments to circumvent these laws. So here's an example from Massachusetts.

So this is a case where Steward Healthcare, who's one of the large systems in my data, a urologist from Steward claims that they used illegal incentives to induce him to do internal referrals. So he claims they would call up his patients and tell them to go to Steward when he referred them out. He claims that when he didn't do enough internal referring, they both

disciplined him and had meetings where they would publicly shame doctors who didn't do enough of this activity. Those are soft incentives. Here's some harder incentives.

They claimed that he -- that Steward withheld a $25,000 bonus for performance issues and eventually terminated him for performance issues, which he denies. Now, lest you think these are the disgruntled ramblings of a deranged, fired doctor, when called for comments, Steward's lawyer called these practices legal and extremely common. Now, that argument hasn't borne out in the courts, but I hope it at least gives you some evidence that the systems believe they can do this and are, in fact, doing this in practice.

Okay. So let me talk about the data a little bit. So we're going to look at two data sources. So the first is the Massachusetts all payer claims database. So those of you who aren't familiar with APCDs, Massachusetts was the first of a number of states to pass a law that required all insurers to submit the universe of their private claims to a centralized database, and we have access to that database. So it's almost all commercial claims for five years.

We have pretty detailed patient demographics.

We have claim line-level data, so I can see things at the procedure level, as well as participation of patients in certain supply-side incentives, which I hopefully will get the chance to talk about a little later.

One thing I should be clear about is that we do not see referrals directly, we only -- because referrals are communication between the patient and the PCP, which are not necessarily reported to the insurer. So we only see where patients eventually sought care.

We link this to a relatively less-used data set called the Massachusetts Provider Directory, which is a database of physician affiliations. So we can match physicians to their practice as well as those practices to a larger organizational hierarchy. So I can see that Ted is an orthopedist in a practice. That practice is in Brigham and Women's Hospital, which is owned by Partners Healthcare.

We observe a single snapshot of this data from December 2014, and we are going to measure -- sort of critically for this talk, we are going to measure vertical integration between the PCP J and Orthopedist K, V-sub-j, k, as binary for whether or not they share any organizational link across the full hierarchy. To give you a sense of what these systems look like,

there's a number of large systems. These are the six biggest.

Massachusetts is relatively horizontally unconcentrated compared to a number of other healthcare markets, but you can see there's a lot of vertical -- there's a lot of vertical concentration in the sense that just these six large systems control two-thirds of PCPs and three-quarters of orthopedists.

Okay, so let me tell you a little bit about orthopedists. So orthopedists are the second highest paid medical specialty. They are second only to plastic surgeons in the U.S., and orthopedics alone comprises 8 percent of medical spending, which multiply that by 18 percent of U.S. GDP gives you that orthopedics alone is nearly 1 ½ percent of U.S. GDP. And this has been a major target of efforts by Medicare to both restrain costs and improve quality.

They have a large volume, and they're -- it's primarily through nonemergent means, so we think they're quite good for thinking about referral patterns compared to cardiology, which a lot of health economists have studied. And they have substantial discretion over the treatment plan for patients who present relatively identically, including surgical options, such as total replacements, to pain management

via either drugs or what are called cortisone injections. We are going to -- our sample of orthopedists is going to include any Massachusetts-based orthopedist who we see performing at least five joint replacements in the full course of our data.

Let me not say too much about them other than the -- after we linked them to the provider directory, you can see on the far right column, the percent of orthopedists who share -- who have a link to at least one primary care doctor is 98 ½ percent. So that's 201 out of 206. They vary a lot in terms of how much surgery they do, sort of unconditionally.

Okay, so -- I've got a little more time. So let me talk a little bit about our sample. So we're going to construct a sample where we want a data set where every observation or sample is a referral handoff. So what we're going to do is we're going to take the orthopedists we identified, we're going to pull all of their claims, and we're going to take the first time any patient who ever saw an orthopedist saw an orthopedist for an office visit, and we're going to dump any patients who it looks like appear in the hospital before they ever show up at an office.

We're going to call that office visit the

moment at which they were referred to the orthopedist, and we're going to look back 12 months from that office visit to retrospectively assign them a primary care doctor based on the primary care doctor they saw the most for office visits in the preceding 12 months. We'll restrict the last three years of our data so we can burn in the first year, so to make sure that these patients never saw an orthopedist for two years when they showed up, and we are going to drop out cases where we either can't find a PCP or where we can't match either of the doctors to our affiliation data.

After that, take a look at the far right column. We end up with about 127,000 patients coming from 4000 primary care doctors to about 200 orthopedists. These patients are primarily women. They are older than the average population in Massachusetts. You can see 96 percent of them have a primary care doctor who shares at least one tie with an orthopedist, and 63 percent of them are sent internally, by which I mean they're referred from a primary care doctor to an integrated orthopedist.

About 20 percent of them receive any surgery -- orthopedic surgery in the first year, and they spend about $14,000 in that year. To give you a benchmark, the average in Massachusetts is 10,000.

All right. So here's how we'll think of patients entering the system, quickly. So we'll think of a patient as incurring some pain in their knee. They decide whether or not to go to a PCP. They go to the PCP. The PCP decides, okay, am I going to refer the patient to an orthopedist or not? Does the pain seem bad enough? We'll then say, okay, conditioned on referring at all, the PCP will -- and patient will choose an orthopedist for the patient to end up at. The patient will end up at that orthopedist, who will choose some sort of treatment course, which will realize some cost outcomes.

I'm going to focus today on the last three parts of this chain. We're not going to be able to say much about patients who never enter the system at all, and we can talk about this in the discussion, but it's hard to think about the extrinsic margin -- hard to measure the extrinsic margin for things like orthopedics, where most people seem to be characterized as having some sort of pain.

Okay. Let me present a really simple model that I hope will motivate both the effects we care about and the identification issues here. So we can think of, again, Patient I as being sent from PCP J to an Orthopedist K. Given that sort of allocation, some

cost outcome, Y-sub-I, j, k is going to be realized after the treatment course, which is going to be a function of stuff about the patient, including what they have, who -- what orthopedist they ended up with and that orthopedist's treatment patterns, whether or not that orthopedist was integrated with the primary care doctor, and some other stuff, Epsilon I, j, k.

Given that structure for the cost function, we can then think of a PCP and patient engaging in some sort of joint decision-making process that generates a structural choice utility function over orthopedists, U sub I, j, k, which we'll think of as a function of, again, stuff about the patient and PCP, the expected cost outcomes, other stuff about the orthopedist, and whether or not that orthopedist was integrated with the patient's primary care doctor.

Given that really general form for costs and demand, we can think of the two objects we care about, these productive efficiencies and the allocative steering effect of integration, as being Eta and T. So we'll think of the productive efficiencies as being the effect of flipping that V sub j, k on and off on cost outcomes while holding all else fixed. Similarly, we can think of the steering effect as flipping that V sub j, k on and off in the utility function, holding all

else fixed.

One thing I should be clear about, which will hopefully proceed to the identification strategy, is holding all else fixed includes a lot more than you think. In the efficiencies, it includes holding the identity of the orthopedist fixed. Similarly, in the steering effect, it includes holding expected costs fixed, which means we have to think of the effective net of the potential efficiencies and responses to those efficiencies.

So given that we can think of how vertical integration affects outcomes to a first order approximation, the effect on demand of J -- of a given J and a given K integrating is going to have some first order positive effect on demand proportional to both the steering effect as well as the extent to which efficiencies are realized and the PCP and patient like to internalize those efficiencies.

Similarly, we can think of the first order effect on cost outcomes as including both the realized efficiencies for inframarginal patients and an effect for marginal patients, which is that patients are going to get re-allocated across orthopedists who induce different costs, and so that, too, is going to have an effect on equilibrium cost outcomes from vertical

integration.

So we can think of the effect of vertical integration, sort of the total effect, as depending on both market structure, sort of these U and Y sub I, j, k's, as well as sort of what we might think of as conduct, these efficiency and steering effects. So we face two big challenges to identification, which hopefully are clear from the model, which are that vertical integration of the two main outcomes includes both the parameters we care about.

So we couldn't just regress these outcomes of vertical integration and expect to get the right thing out of it. We are going to solve that by using a two-step process where we first estimate the efficiencies as the spending effect of VI conditional on the orthopedists you see, and then we estimate the steering effect as the preference for integrated orthopedists conditional on the efficiencies we estimate.

Second, given that nearly everyone is integrated, it's hard for us to think about a control group of unintegrated doctors to think of the counterfactual where we might think of breaking everyone apart, and our solution is we are going to use variation on the integration of pairs, so everyone is

integrated with someone, but not someone is integrated with everyone.

We are going to have to use some no-sorting assumptions to allow that to work, but given those assumptions, we can use the variation of who's integrated with who to identify all of our effects. What's nice here is we -- we are going to get the effect of pairwise integration, which is not necessarily affiliation with a given large system.

Okay, I'm running a little short on time, so let me go faster. So we are going to start by estimating a -- actually, I should run through that just a bit. We are going to estimate a cost function for orthopedic treatment. Given those estimates, we're going to estimate a demand function for orthopedists, and we're going to use those parameter estimates to simulate counterfactuals where we break apart the vertical ties we see existing.

So we are going to model Y sub I, which is going to be one year all cost spending after that first -- after that index visit as a log linear function of orthopedist "fix fix gamma sub k," a term for whether or not the orthopedist -- I'm sorry, the patient's orthopedist and primary care doctor were integrated, as well as a rich set of observables we see

about the patient. We'll think of that gamma sub k as the risk-adjusted cost of seeing a given orthopedist, k, and that Eta term, which is the coefficient of V sub j, k, as a measure of efficiencies.

I'll be clear that our identification here really relies on no sorting of patients across orthopedists on unobservables, and so Eta is going to be identified by the within orthopedist variation and the integration status of the primary care doctor that the patient comes from.

I know no certain observables is a hard assumption to buy. We do have evidence that there's limited sorting on observables, and in a sort of standard Altonji and Taber test, it looks like we will actually underestimate the efficiencies we eventually estimate.

So here are our estimates. So the standard deviation of those gamma sub k's, which you can think of as sort of the dispersion -- roughly the dispersion in productivity of orthopedists, is 0.3. That's in log points, so you can think of that as going from the average orthopedist to one who is one standard deviation more expensive, results in about a 30 percent increase in expected cost in the first year. That's about $4,000.

That's substantial. The efficiencies, too, we estimate as being relatively substantial. So we estimate them as a mean effect. They result in about a 6 percent reduction in expected spending outcomes in that first year. That's -- I always forget the math on this. I think it's about $700. And those effects are really heterogenous across the system.

So Partners, who's considered the high-cost operator in Massachusetts who employs a lot of your favorite Harvard Med School professors, they seem to realize lower efficiencies, whereas Atrius, who's a sort of smaller group, seems to realize extremely large joint efficiencies.

This is a -- to give you a sense of how these -- this is a patient-weighted histogram of these orthopedist fixed effects. Zero, again, is the average orthopedist, and you can see even though there's substantial mass at orthopedists who are 20 percent less costly, there's substantial mass at those who are 20, 40, even 60 and 80 percent more expensive than the average orthopedist -- sorry, incur 80 percent more costs than the average orthopedist.

Normally I would have a clicker and I would show you something, but let me claim to you -- and you can read the paper to see this -- to see the graph

here -- that you were worried that this is quality, this is not quality. If you scatter my cost estimates against measure -- against ProPublica's surgeon scorecard measures of hip and knee complication rates, you get a slope of almost exactly zero. So these cost differences are not representing -- don't seem to represent differences in outcomes across patients, and the efficiencies seem to largely come from things like imaging, where we think there's a lot of waste, rather from directly provided orthopedist services.

Okay, so given those cost estimates, we can now estimate demand function for orthopedists. We're going to estimate this as a logit model, where utility is a linear function of orthopedists' costs, which we'll represent as our cost estimates, plus the efficiencies when they're relevant, this vertical integration term, and a couple other things about the orthopedists, including their propensity of surgery, distance, quality, and dummies for what large system they're a part of.

We also allow sensitivity to cost and other things that depend on patient observables, as well as whether or not their primary care doctor was subject to incentives that made the primary care doctor more likely to seek out a lower cost orthopedist.

The identification here is going to rely on the idea that integration only affects choice utility at orthopedists you integrate with. What that means we can do is we can use, say, Atrius PCPs as preferences for Partners versus Steward -- those are two of the big systems -- as the counterfactual for Partners' PCPs, as preferences for Partners versus Steward if they were unintegrated, and we'll use sort of every pairwise version of that to identify our steering parameter.

Here are their own utilities, so they are not directly interpretable, so let me interpret them for you. So, sorry, the first column is the coefficient on orthopedists' costs, and the right column -- and the right column is the sort of steering parameter. So let me tell you that that cost sensitivity is essentially zero. It's equivalent to a one-deviation change in the sort of standard logit idiosyncratic preferences, is equivalent to, in utility terms, a 200 sigma change in orthopedists' costliness, which is well outside the domain of our estimate -- of our data. So you should think of these -- this demand as being essentially cost-insensitive.

We also see that T -- the steering parameter is surprisingly not increasing potential system rents. So Partners really has the most to gain from steering

patients internally, because they're very expensive, so they can command high revenues, and they own expensive hospitals. Despite this, the system that seems to be doing the most steering is Atrius, who actually has the least rents, because they don't own hospitals, and so if you thought sort of the story of how the steering occurs is some sort of internal rent-sharing, that could happen roughly, but it doesn't seem to, at least in the aggregate, when we compare these systems to each other.

Okay, so given those estimates, my last few minutes, let me show you some results from counterfactual simulations where we remove vertical ties from our data. So we'll simulate breaking -- we'll simulate removing the efficiencies and we'll simulate breaking the ties altogether, so that will be setting those Etas to zero and then setting all those V sub j, k's to zero. We'll measure the effects on both these internal referral rates as well as the expected cost outcomes. We'll do that by -- when we change the parameters, we'll recompute market shares, and then we'll compute expected cost outcomes from the stages of our model.

And I should be clear, we're abstracting from GE price responses, which I think could be potentially

substantial, and I'm happy to talk more about what those could look like.

So here's the rate of internal referring. This is the percentage of patients who are going to be sent from a PCP to an integrated orthopedist. We can see in the status quo, our counterfactuals give us roughly what the data is, which is the rate is about 63 percent. Removing the efficiencies takes that down by only about a percentage point and a half, whereas breaking the ties altogether cuts that by about 60 percent. Breaking the ties altogether re-allocates about one-third of patients. So these ties are really, really influencing where patients are allocated.

So given that large allocation effect, you might think that this sort of steering is bad for costs, that it might increase costs. We break -- when we break all the vertical ties, we find the opposite. So breaking all the vertical ties increases costs -- expected costs by 6 percent. Where does that come from? Well, a lot of that comes from losing the efficiencies that exist given the high rate of internal dealing, but despite that, when we actually take away the efficiencies, taking away the steering effects conditional on -- even after taking away the efficiencies still increases costs slightly, and you

can get a hint to why that is when you look at the heterogeneity across the system. So removing Partners lowers costs, but removing Atrius or Steward, who are relatively lower cost, increases costs.

And the story running in the background is the fact -- the reason for this is that we have no demand sensitivity to cost, so in the absence of these steering incentives, patients are not sorting to lower cost orthopedists. They're sorting orthogonally to orthopedists' costs. And so the low average cost systems, like Steward and Atrius, seem to be -- their steering seems to be providing actually positive effects on patients' expected cost outcomes.

Now, if you want to take that plus 1 percent effect down to zero, you'd need that beta, the cost sensitivity parameter, to be about 0.8 higher. That's about 40 times higher. That's a really high bar given that, in my own past work and other work, we've shown that demand-side policies don't really seem to get you any of the way there. In estimates from this paper, we show that supply-side incentives get you part of the way there, but only about two-thirds. Let me skip to exactly how far they get you. Yeah, they get the steering effect down, but actually they raise the detrimental effect of removing efficiencies.

Okay, in my negative one minute, let me say something very quickly about antitrust given it's the FTC. So we know the market study sort of has a nice summary of this in a review paper from this year. Ex ante evaluation of vertical mergers is going to require us to estimate both the horizontal market structure and the vertical conduct. What we're finding suggests that there's substantial heterogeneity across firms even within a given market in conduct, and so there's no real -- we estimate there's no real singular impact of vertical integration, and so data from the experience of other organizations may be completely uninformative about the experiences of future vertical mergers. So we really need better models and, frankly, data on within-firm conduct in healthcare to think about, you know, evaluating potential effects of vertical mergers.

Let me just say to sum up, we -- you know, given that we find these large effects in vertical integration on healthcare industry productivity and given the lack of cost sensitivity, we should really worry about dynamics here. So without cost sensitivity, we might not get the nice Schumpeterian dynamics we expect in lots of other industries that are a little more functional than healthcare, and so policy towards mergers which affect the dynamics of

integration may really be influential for healthcare industry productivity going forward.

Okay. Thanks for your time.

(Applause.)

MR. VIOLETTE: Thanks. Now we have Jose Cuesta from Stanford to discuss.

MR. CUESTA: Okay. Well, thanks for inviting. It was a fantastic paper to read and I think a very interesting case study of VI in the health market. So just to put things in context, this is essentially the -- kind of the big picture of health markets, so we have downstream consumers demanding insurance from different insurers. They pay a premium. They get the -- you know, some plan. They get a network. Once they get a network, they go to the hospitals, right? And then within hospitals, we have a -- kind of an ecosystem of its own, right?

So what -- what Zarek is doing is actually constraining the analysis of vertical integration to what happens within that organization, right? So, in particular, even when I think of how VI, vertical incentives, kind of like may affect the outcomes within the firm, there's two key forces that they highlight quite a bit. First, efficiencies, so incentives to reduce cost. Second, steering, so incentives to move

patients around towards integrated physicians, right?

So what they do is to focus on a very compelling case study of how PCPs refer patients to orthopedists, right? And (indiscernible) is a patient gets to a PCP, and that's a -- you know, outside of the model, and then when they get there, the PCP receives them and then refers that patient out to given orthopedist, right, okay, which might or may not be integrated, so it might be a referral within the organization or through -- or to a different organization, right? And then the orthopedist receives the patient, treats him, and then the cost, Y, is realized, okay? That's the framework that they develop, okay?

So what they do in practice is to use data on choices. So this is a choice of -- the joint choice of the PCP and the patient of where to go in the upstream, right, and cost, so realized cost to do three things, right? So first -- two things, essentially. First they estimate this efficiency on steering effects, and they do it using a very nice framework, I think. And secondly, they do, what happens if we had done VI, that is, if we break all those -- if we essentially, like -- we remove the constraints that the organization may impose in the relationship between these two levels of

physicians, right? We do that keeping all the rest of the market fixed, right? So that's one thing, and they acknowledge it in the paper and their representation as well.

The main finding is going to be, actually, that in this setting, VI reduces cost by 6 percent, and that most of it comes from the steering within local hospitals combined with this lack of cost sensitivity, okay? So that's the paper. So I think this is going to be essentially repeating one of Zarek's slides, but I think it's a fantastic slide, actually, because it makes very clear what the framework is.

So the model can be summarized in these two patients. One -- and I am going -- the way I think about this, this is a selection model, right? So first we have the outcome equation, which is cost, and it's going to be a function of -- well, the identity of the orthopedist and VI and some other attributes, and then we have the selection equation, which is in a -- you know, we have a patient, and we send it to some -- to some orthopedist, right? And that's going to be, again, a function of patient attributes, of expected cost, and then the orthopedist's attributes. So this is essentially a choice model and, again, integration.

They write down these two very nice expressions

for what the productive efficiencies are, Eta, which is one of the parameters of interest, just -- you know, the change in cost if you shut down vertical integration and the steering effect, okay, which is T, which is going to be the difference, if you want, in the choice probability of the PCP towards a given -- a given orthopedist when you shut down vertical integration, okay?

The challenge here is that Eta and vertical integration enter the steering effect, and, therefore, they need to seek some strategy to tear those two apart, okay? What they do is to develop this two-step strategy. First, they are going to condition on the orthopedist and estimate the cost function, that's kind of -- some kind of a statistical model of cost, right, shutting down any selection on unobservables, and then in the second stage, they go back to the first stage and estimate steering effects given their estimates of Eta, okay? So that's I think the framework.

So I have a few comments. Hopefully they are going to be constructive. Hopefully they are not deceptive after today's talk. So, first, VI, of course. This is a big theoretical literature, and I know there's increasingly more and more work trying to do empirical work in the area. So the paper focuses on

this tradeoff between efficiencies and foreclosure, which is, you know, substantial in this literature. Importantly, so you read all these papers, they always start, you know, the literature is not settled on what are the effects of VI. It's more likely the case that the result is going to be industry-specific and is going to depend on the specifics of each industry, right? In that sense, the answer to that kind of claim is maybe we need more case studies, and this is, indeed, a very nice case study of a particular and compelling industry, which is healthcare, okay?

So my comment here is -- okay, so, now, given that we are going to a case study of healthcare, can we learn more about what VI does in this industry, right? So we have first the estimate of Eta. Most of the attention in the paper and to some extent in the presentation is focused on the estimate of gamma, which are the physician fixed effects, right, heterogeneic – or the orthopedist, actually, right? I could argue that it would be more interesting to think, okay, so what determines Eta and is there -- what's underlying that, right? What are the practices that change across integrated and unintegrated orthopedists that are able to reduce costs, right? And they do have more granular cost data that they could use to decompose that --

those effects, and then to try to link them to this literature on physician behavior that is focused on moral hazard, selection, information-sharing, financial incentives, right? Try to go back to that literature.

Second, there's a -- on the orthopedist side, there's a thing -- there's a concept of foreclosure. That is, maybe they don't get referrals because they're not integrated, right? Can they react in some way? Can they do investment or change attributes such that they're able to become attractive even though they're not integrated? That would be interesting as well.

And finally there's the issue of within the organization what's going on, right? What are their financial incentives? What are the payment structures that are, like, driving this behavior? And there's the literature on exclusive dealing, which looks very much like the story that they are trying to tell, which would be interesting to speak to, okay?

Okay, so I have three minutes. So I am going to skip this one. I want to say one thing here that I think is important. So as I said before, one way to think about this problem is as a selection model. There is some recent work in the area of either estimating returns to college, which is essentially the same interesting problem, right? You apply to college.

You go to one. There's a substantial selection of unobservables in that first stage, and then you get returns in the second stage of the model. This is the same.

Similarly, the paper by Peter Hall estimating quality, that's the same problem, right? Both of those papers are using distance as shifters, right? And with that, you can actually kind of identify a model that allows for selection of unobservables, right, which is something of the paper's -- sorry?

UNIDENTIFIED SPEAKER: No, I'll tell you.

MR. CUESTA: Okay, good.

Okay, so I don't know, that seems like an avenue to kind of, like, relax the assumption of selection of unobservables that could be helpful, okay? And then, again, there's a few comments about the comparison group, but that's -- that's fine.

Okay, then I think this is something interesting. So in the paper, the way in which the utility function of the PCP or of the first problem as written is as follows, is there's a same choice utility. There's going to be a weighted function of the utility from the -- of the patient, which is that Thi V, and then there's going to be the utility of the physician, which is his compensation, say W, right? So

we're thinking of these pairs, solving this joint problem, and then going to a physician, right?

In the setting, there's actually quite good variation to identify the utility of the physician, because they have a variation in payment incentives, right? However, identifying alignment or altruism, how it's called in the literature, is usually hard, right? So because it's -- it can be separately identified from consumer preferences, right? So in practice what they do is to estimate in practice our reduced-form model, which puts all of it -- all of this somewhat together, right?

And then a limitation of that is that it's hard to do counterfactuals related to some financial incentives, and, importantly, it limits the extent to which we can do welfare analysis and consider what happens on the patient side of the market, right? And that's fine. I mean, the paper is pitched as a paper on productivity and cost, right?

So how to improve? I wonder if there's, like, some segment in which the consumers are actively choosing, right? Could we then go to that segment, recover consumer preferences in that segment, come back to the original (indiscernible), and then recover all the parameters and do counterfactuals along the line,

for example, of, you know, different constructs, do welfare analysis, and think also about the misallocation of patients across different orthopedists that vary in welfare-relevant attributes, right?

And we know from previous research, for example, from the AI -- our paper by Martin Gaynor on free choice in the UK, that actually choice can improve welfare, right?

Okay. Then there's this comment and then I'll wrap up. So a lot of the results hinge on the finding of lack of cost sensitivity, right? They're estimated essentially at zero on cost sensitivity, and I understand there's, like, very good -- at least one variable by a produced paper showing that there's limited cost sensitivity.

However, there's like one very salient choice, right? So maybe we could think of a -- you know, maybe there's something wrong and we're not capturing that cost sensitivity well in our estimates. I have one suspicion of what could be going on, and this is not clear in the paper, and maybe you can help us out afterwards, but in the market, there's like 200 orthopedists. The average PCP refers to nine.

So I wonder, do they actually consider all those orthopedists in their choice set? If they're not

and they are specifying the choices as all the orthopedists in the market, then we might actually be getting misleading estimates of our preferences, right? And in particular, we might be underestimating cost sensitivity, right?

So can we use something better? Here our suggestion is maybe estimate -- kind of like try to integrate the choice sets, right? In particular, you could think of, you know, some PCPs only refer within their VI chain. Some PCPs consider all the other ones, and the weights on those two choice sets would inform us about how they actually behave. I don't think that the average result will change, but I think the interpretation will be richer, and also we can think of heterogeneity, okay?

Okay, so two -- just to close -- yeah, so, you know, it's not extensive, but thinking broader about the health industry, so what happens downstream? Do any of these efficiencies or lack of, when you ban VI, actually pass through to hospital prices and then to premiums, how that changes welfare, how that changes choice downstream? Those are obviously interesting questions, beyond the scope of the paper, but at least a discussion of the relevance and the potential magnitude of them would be super-interesting.

And then -- okay, so this is a -- so the counterfactuals essentially are quantifying the effect of VI, right? Could we think more about policy or actual regulation, say one thing is -- so there's like a lot of the stories about information-sharing, right? We could have, perhaps, like national or like more widespread electronic records? That's something that can be seen related within the framework itself, right, and that could, you know, make the efficiencies of a VI shared beyond the boundaries of the firm.

Okay, yeah, so exciting paper, exciting agenda. I really enjoyed it. Yes, thank you.

(Applause.)

MR. BROT-GOLDBERG: Let me just say, that was a great discussion, and let me just say three things about things Ignacio said. So on distances and instrument, it doesn't work unfortunately, and the reason it -- it's -- it does work in hospitals, but it doesn't work in physicians, because everyone's densely packed in Boston. So you can't -- most -- a bunch of our physicians are literally in the same location.

On welfare, it's hard to think about welfare here even if we were successfully able to separate out the patients versus the PCPs. Should we really think about the patients' preferences as welfare? Given my

other paper, where we see extensive intensive price margins that look really different, my answer is no, but I'm happy to talk more about that.

And then -- oh, and the market definition. Yeah, this was a very peculiar thing. I was hoping to do it by HRR. It turns out the market -- the relevant market is -- looks like the whole state in orthopedists, because half of people -- half of my patients in Western Mass drive into Boston, and so it was really hard to think about who exactly is in the choice set, although it's a great suggestion to think about, like, technical ways to try to think about what the choice set might be.

Cool, all right.

AUDIENCE: (Off mic) variation in referral patterns in practices make it possible (off mic) issue referred to -- thank you -- referring to patient cost share, overall cost, is the first thought.

Then the second -- and you might want to look at California, I think, and Alpers (phonetic) had some other projects with incentives for going to centers of excellence or physicians that are more efficient and have better outcomes. It may be interesting for -- you know, if you can get some data from California. It's just a suggestion.

The third thing about anti-kickback statutes, I'm not sure you're going to find this data, but the Federal Employee Health Benefits Program is exempt from that anti-kickback statute, so I wonder if you would find more --

MR. BROT-GOLDBERG: Sorry, what -- who's exempt?

AUDIENCE: The Federal Employee Health Benefits Program.

MR. BROT-GOLDBERG: Ah, I can't identify -- I don't think I can identify them in my data, so that would be very difficult.

AUDIENCE: No. Right now, no, I guess --

MR. BROT-GOLDBERG: And it doesn't seem to be super-relevant anyway given that everyone -- you can essentially engineer a kickback scheme within the firm.

AUDIENCE: Yes.

MR. BROT-GOLDBERG: -- on incentives. So what we used was all costs -- all allowed expenditures, so that includes both what the patient and the insurer pay. The reason we think that matters is because someone has to pay those costs eventually, and the patient often has to pay them through pass-through premiums anyway. So we think it's relevant.

If you think that patients are not sensitive

because they're not bearing the cost, I think, you know, I would accept that, but I think my other paper suggests that that's probably not what's going on. And something I didn't really describe but is in the paper, we have sort of in Massachusetts major insurers put incentives on PCPs, made them bear incentives to refer to -- to refer to -- to keep costs down, and that you can see had large effects on cost sensitivity.

So, yeah, so all the counterfactuals are zeroing those incentives out, and the table I skipped over puts those in, and it changes the effects by a little bit but not a lot.

Are we good? I thought I saw someone. Okay, we're good.

(Applause.)

MR. VIOLETTE: Next we have Claudia Robles-Garcia from Stanford presenting on "Competition and incentive in mortgage markets: The role of brokers."

MS. ROBLES-GARCIA: Okay. So thank you very much for including the paper in the program, and thank you for all of you for staying up so late. I know I'm keeping you from the drinks, so I'll try to make it as painless as possible.

Before I dig into mortgage markets and the role of brokers, let me give you a big picture of the type

of setting and market structure I'm going to be thinking about today. Okay, so when you think about transactions nowadays, when consumers want to purchase a good or a service, more often than not, they do it via an intermediary, and intermediaries sometimes act in the context of expert advisors. So if you think of an actual product, when consumers want to buy a mortgage or a credit card, they often go to a dealer or a broker to get some sort of advice.

Now, you might be worried that the way these brokers or dealers get paid is going to affect their incentives and, finally, their recommendations to consumers. There's been a substantial policy debate not only for brokers or dealers but also for physicians or some sort of intermediaries that are dealing with consumers which might be less informed than them. So I'm going to try to contribute to this policy debate on how we compensate expert advisors by looking at mortgage markets and the role of mortgage brokers.

So mortgage brokers are essentially an intermediary between consumers and banks, and in the context of the UK, which is the market that I am going to be talking to you about, they account for 50 percent of all mortgage originations. If you look at other markets, such as Canada or the U.S., they have a

significant market share as well. So what exactly do these mortgage brokers do that makes them so popular?

Well, consider the mortgage market. We have consumers, borrowers. They found a house that they like. They need a mortgage to actually buy it. And we have banks offering different types of mortgage products. Consumers can follow a more traditional approach and go directly to the bank. So you could think of these as walk into your nearby branch or buy the mortgage online.

Alternatively, consumers can also desire to hire a broker. So they are going to pay a fixed fee to a broker, and these brokers, more often than not, belong to large broker companies. So these broker companies are going to essentially provide two types of services. First, they are going to provide advice. So they are going to give some sort of recommendation to the consumer and explain either what a mortgage is or basically which products are available in the market. And the second service they're going to provide is in terms of all the paperwork. So for those of you who have a mortgage, you probably know that it's a painful process, and the broker is essentially going to speed up the process.

Now, this fee that brokers get paid by the

consumer is not their only source of revenue. More often than now, brokers also get a commission payment from the banks. Now, these commission payments are usually a percentage of the loan and are broker-banker specific. So the way they work is that each bank and each broker company negotiate these commission payments, and they're very heterogenous, both in the U.S. and in the UK.

Now, you might be worried that because brokers get paid by the banks, these might potentially lead to an agency problem between the broker and the consumer. So if you think that brokers steer consumers to products that have a higher commission and these products end up being more expensive for the consumer, this can potentially be detrimental for borrowers.

Now, the point that I want to make today in the talk is very simple. On the one hand, brokers may have a negative effect on consumers if there's evidence of this agency problem, but it also may have a positive effect on consumers by increasing efficiency or competition between the banks in this market. So mortgage markets are very concentrated markets, and if brokers in some way increase competition among the banks, if this brings interest rates down, this can actually lead to sort of general equilibrium effects

that could, at the end of the day, be beneficial for consumers.

So given that there's a potential tradeoff, what are regulators doing about this? Well, if you look at many markets, not only mortgages but also credit cards and physicians, you see that in many markets, including the U.S., regulators have decided to ban all payments between providers and intermediaries; in the case of mortgages, between banks and brokers.

The point I want to make, very simple. If you ban commissions between providers and brokers, this is going to reduce the agency problem because it's going to align assymetries between consumers and brokers. However, this may also have unintended consequences when it comes to competition among the banks, as well as to efficiency in this market, and at the end of the day, what is the overall effect on consumers is going to depend on which of these two forces dominates in equilibrium. And so theoretically I'm going to show you today that it's ambiguous, and we're going to need some sort of empirical evidence to see which of these two forces dominates in a given market.

So what do I do in the paper? So as I mentioned before, I am going to be looking at the UK, and I am going to use this amazing data set that

captures all mortgage originations in the UK at a very detailed level. So whenever you deal with intermediaries -- there's a few in the audience that have already worked with these types of agents -- you know that it's very hard to get information on the payments they get from providers. So it's very hard to get micro-level data on these commissions. This data set is going to allow me to observe, for every mortgage, which broker originated the mortgage and every single payment the broker received, both from the consumer and from the bank.

So this level of detailed data is going to allow me to estimate a supply and demand model which essentially is going to capture this tradeoff that I just mentioned, and I am going to use this model to answer three research questions. The first thing that I want to check is what is the distortion that these commissions are causing in the broker's choice? So first I want to see whether brokers are reacting to changes in commissions and whether this essentially creates an agency problem or not.

The second thing that I want to check is whether brokers have a positive effect on consumers, so whether in any way brokers are increasing efficiency, maybe by reducing borrowers' search costs, by reducing

bank marginal costs, or whether they're increasing competition among the banks. I'm going to find evidence that there is an agency problem, so negative effect for consumers, but there's also a positive effect on competition and efficiency. So at the end of the day, I want to understand, if we were to regulate the way these brokers get paid, so if we were to impose a ban or a cap on these commissions, which of these two forces is going to dominate and what is going to be the ultimate effect on consumers?

Okay, so essentially there's been a huge literature trying to look at intermediaries, but when you zone in on the incentives of these intermediaries, there's very limited empirical work being done, and one of the main reasons has been data limitations. As I mentioned before, it's hard to get data on these payments, and most of the approaches so far have been focused on the demand side, on the relationship between the borrower and the broker, between the consumer and the intermediary. If you only look at this part of the market, there's no tradeoff. If you ban commissions, the agency problem shuts down. However, if you look at the general equilibrium, at the relationship between the providers and intermediaries, things may start looking different. So that's where my paper comes in.

Okay, so let me tell you a bit more about this amazing data set on the UK mortgage market. So as mentioned before, I'm going to observe every mortgage origination for about a year and a half, 2015 to 2016, for the UK mortgage market. For every mortgage originated, I am going to observe a very rich set of mortgage characteristics, a very rich set of borrower characteristics, and, most importantly for this paper, I'm going to observe which broker originated the mortgage, if it was intermediated, and every single payment the broker received.

On top of that, I am going to observe every single contract that the bank and the broker signed during this time period. This is going to be very important because when I get to the model, I am going to endogenize this relationship between the bank and the broker, because not all brokers are going to deal with all banks, and not all banks are going to deal with all brokers.

And, finally, if you think about the outside option of the bank when negotiating with a broker, you could think of this as being the branch, the direct sales, because as a consumer, you can always go directly to the bank. So I'm going to observe also every single branch location during this time period.

Okay. So I'm talking about the UK mortgage market. You don't really need to know anything about it. So let me give you this crash course on UK mortgages. So the first thing you need to know about the UK mortgage market is that unlike the U.S., there is very limited individually specific pricing. So what do I mean by this? In the U.S., two individuals going into the bank may face different rates for the same exact mortgage, depending, for example, on their credit score. If I have a better credit score, I might get a better rate for exactly the same product. That is not the case in the UK.

In the UK, two individuals with the same credit -- with -- getting the same mortgage will get exactly the same rate conditional on approval, okay? So the story that I am going to be telling you today is not about the broker negotiating a better rate on my behalf. That's not what's going to happen here. What's going to happen is the broker finding me a better product from another bank that I might not have been aware of. So that's the mechanism I'm going to have in mind.

The next thing you need to know about the UK mortgage market is that it's very concentrated upstream. We have what we call the Big Six lenders

that account for more than 75 percent of all originations, but what's going to be very important in my survey is that in the recent years, we've seen entry of what we call challenger banks. So the challenger banks have a very different model from the Big Six. The Big Six are traditional banks. They have a lot of branches. They hold all the mortgages on their balance sheets, and they rely heavily on advertisement.

The new banks are a bit different. They don't have many branches. They don't incur high costs in advertisement, and the way they're going to introduce their products in the market is by using brokers. So what I am going to say in the paper is that it is these new banks, the ones that are offering a higher commission to brokers. So that's the way they're going to introduce their products in the market.

And the last thing you need to know about the UK mortgage market is that it is also very concentrated at the broker level. So when you think of a broker, don't think about this one-person firm industry. No, it -- these are big companies. So you have the large 20 broker companies accounting for more than 66 -- sorry, 65 percent of broker sales, but once you zoom in at the local level, at most you see four or five brokers. So at the local level, these broker companies

are going to have significant market power. So you could think of this industry as a bilateral oligopoly.

So how does this industry work? As I mentioned before, almost 50 percent of all borrowers use a broker. For this paper, I'm going to be focusing on first-time buyers, and in this submarket, 70 percent of them use a broker. How do brokers get paid? Most of the income a broker gets, they receive it from the banks. So most of the time, consumers pay zero for the broker services, and they're going to -- they're completely aware that most of the income the broker gets is going to come from the bank.

Now, how much they're going to receive from the bank, it varies a lot, so there's massive heterogeneity both through the same broker, across banks, but also from the same bank, across brokers. So you could have one bank offering a very high commission to a broker, a very small commission to another broker. And there's also a lot of heterogeneity on how many brokers does a broker -- a bank have and how many banks does a -- how many banks a broker has and how many brokers the bank operates with.

Okay, so before I move on to the model, let me give you a bit of a taste on what the data looks like. There's a lot more information in the paper, but let me

give you at least a high-level correlation. So the first correlation that is very consistent across the data is that brokers tend to sell products with higher commissions. This is true both in the cross-section as well as in the time series. So for the same product, same set of borrowers, if the commission goes up, the broker is more likely to sell that product.

The second correlation that is very consistent in the data is that whenever there's entry of brokers in a given county, competition among the banks goes up. So it seems that whenever a new broker comes into a county, suddenly these new banks start getting larger market share, and the reason that -- what's driving this increasing market share of the new banks is that when consumers go directly to a bank, more often than not, they tend to go to the nearest branch. New banks don't have branches. They find it very hard to access consumers.

So essentially I'm going to develop a model and try to see which of these two tradeoffs is going to dominate. Okay, so it's going to be a static equilibrium model, and it's going to have three key players. We're going to have households, borrowers, banks and brokers, and each of these players is going to face sequential decisions. So the way it's going to

work is the following.

First, every broker and every bank are going to meet, and they're going to decide whether they do business with each other. So if they reach an agreement, the broker can sell the bank's products at a given commission rate. If they don't reach an agreement, the broker cannot sell the bank's products. So this is essentially a network formation.

Now, once all these negotiations between banks and brokers are over, there's going to be -- each broker is going to have a portfolio of banks with whom they can do business. Once we have all this network in place, each bank is going to decide how to price their product. So they're going to get interest rates. So this is the supply side of the model.

What about demand? What are consumers choosing? Well, consumers are going to face two sequential decisions as well. First. They're going to decide whether they go to the broker or they go directly to the banks, and conditional on this decision, they're going to choose a product in their available choice set, okay? This is a sequential model. We solve it backwards. So I will try to be not very painful on the equations.

So first consumers need to decide whether they

go to the broker or not. I'm going to assume that each consumer has a search cost. If they go directly to the banks, they're going to have to pay the search cost. If they go it the broker, they are going to have to pay a fee, but the broker -- but they don't have to pay the search cost because the broker is going to help them get a mortgage. So at the end of the day, what consumers are deciding is the following.

On the one hand, I need to pay the search cost if I go directly to the banks, and I'm going to get some expected utility from the banks. On the other hand, I can pay the fee to the broker, I will match -- I will be matched to a given broker, and I will get some expected utility from going to that broker, okay? Binary choice, depending on search cost versus the utility I will get from the broker.

Now, conditional on going directly to the banks, consumers are going to choose the product that maximizes their indirect utility. What do consumers care about? They're going to care about the price, the interest rate, some observed and unobserved product characteristics of the mortgage, such as the loan to value, the initial period, and so on, and I'm also going to allow consumers to care about nearby branches, and this is something that I see in the data that

consumers more often than not go to the nearest branch when they go direct.

Consumers going to the broker are going to face a different problem, because now we have two agents as opposed to one taking the decision. I am going to assume that the broker and the borrower maximize the joint utility, and the joint utility is essentially a weighted average between what the borrower wants and what the broker wants, okay? So the broker is going to care about how much the bank pays the broker, so they are going to care about the commission rate, and they are also going to care about some sort of cost, which I'm going to have to estimate.

Now, if you look at the parameter theta, this is essentially how much can the broker extract from the consumer. So if we live in a world of benevolent brokers, theta should be zero. If we live in a world of perfectly naive consumers and no dynamic incentives, theta should be one. I am going to estimate theta and going to find that it's about 0.4. So brokers are not maximizing borrowers' utility, but they're not giving them the worst product either. So there seems to be some sort of bargaining taking place between the two.

Let me skip this slide. This is essentially saying that when you move to the supply side and banks

need to choose interest rates, they're going to maximize expected profits, and I'm going to use the first of their conditions to back out marginal costs. So this is an IO audience, so I feel like I can skip this slide.

Let me tell you about the last stage of the model before I show you the results. I'm going to assume that at the beginning of each period, every bank and every broker meet, and they're going to negotiate whether to form an agreement. I'm going to assume that these negotiations are Nash in Nash. What do I mean by this? Very simple. I am going to assume that all negotiations are going to take place at the same time, and once they are over, I don't allow for renegotiation, okay? So this study is very common in bilateral oligopolies, and it has its limitations, but for computational purposes, it's what we have.

So what is each bank and each broker maximizing? They're going to maximize their joint surplus. However, this is a bargaining, right? So you cannot force them to reach an agreement, and at any time either party can walk away. So I am going to have two participation constraints.

Now, the participation constraint of the bank, of the lender, is going to tell me the maximum

commission the bank is willing to pay the broker. The participation constraint of the broker is going to tell me the minimum commission the broker is willing to accept. Sometimes what the bank is willing to offer is less than what the broker is willing to accept. This implies that the link is going to break and they're not going to reach an agreement, okay?

So the fact that I have these participation constraints is going to allow for this endogenous information of the network, which is going to be -- which is going to allow me to change the network once I get to the counterfactuals.

Okay, so since I'm running out of time, let me skip the identification. I'm happy to talk about it offline. I'm sure -- I'm sure Jean-Francois is going to comment on it, too. So let me just show you the results in the last five minutes.

Okay, so what I do find? I find that, as usual, I find elasticities that are common in the literature for mortgages, higher interest rates, less willingness to pay, but what is particular in my study is two components. One is that consumers going direct have a strong preference for nearby branches. This is going to affect the small players that don't have branches because it's very hard for them to access

consumers directly. So the only way we can access them is through the broker.

I also find that the theta parameter, as I mentioned before, which is what I called the broker distortion, but it's essentially how much surplus can the broker extract from the borrower, is about 0.4. I can plug this parameter for every single broker company in my sample. So each dot is a broker company. I see that on average they're all different from zero, so I can reject the possibility of a benevolent broker; however, I cannot reject that they are different from each other.

So it seems that there is some sort of mechanism -- which I am not going to take a stand on in this paper -- that there's a common incentive throughout these brokers -- and hopefully that will be the next paper, so happy to do advertisement later -- but you cannot really separate them that much.

What else do I find? I find that in this industry, search costs are very high. They account for 20 percent of consumer surplus. So in terms of efficiency gains, the fact that brokers are there is going to increase efficiency by reducing search costs for consumers.

I also find that brokers have a more efficient

technology when originating mortgages than banks themselves. So if you look at marginal cost, marginal cost for banks of originating the mortgage with a broker is actually lower than directly through the branches.

I also find that when you look at the cost of the broker, brokers also prefer to originate mortgage with the big banks because the cost of originating a mortgage with the new banks is actually pretty high.

And I also find that if you don't take into account this bargaining in this endogenous network, you're missing on a lot. So given these parameters and these estimates, let me show you the main counterfactual.

So the main counterfactual is going to be the following: Imagine we restrict the commissions between the bank and the broker. So you can either implement a cap up to all the way to a ban. So let me explain the graph in the -- in the screen. On the X axis, I have the cap, right? So you have a cap equal to zero, it's equivalent to a ban. If you have a cap larger than 0.9 percent of the loan, it's as if there was no restriction.

On the Y axis, I'm going to have the change in consumer surplus, because at the end of the day, as a

regulator, you want to know whether you can do better than the current allocation. What I find is that, as you increase the cap, consumer surplus initially goes up, it hits an inflection point, and then it decreases at a faster rate.

So how do I interpret this result? I like to think of this result as essentially a fight between broker market power and bank market power. Consumers going directly to banks are exposed to banks' market power; in particular, since consumers going direct go to the nearest branch, new banks don't have many branches, so there seems to be like this holdup problem between big banks and consumers when they go direct. So that's the banks' market power.

When you go to the broker, the broker also is able to extract surplus from you, and it is this theta parameter that I told you about. So you could think of these as essentially, like, one of the two evils -- don't quote me on this -- like it's essentially this idea that there's two agents, there's two oligopolies, each has market power. Which of them dominates is going to depend on the (off mic).

So let me give you a bit of intuition on the ban on commissions in my last minute. So imagine we impose a ban on commissions. So now consumers are the

only ones paying for the broker. So this must lead to a fall in the agency problem, right, because now the broker is not influenced by the banks. Therefore, they -- suddenly the incentives between the consumer and the broker are aligned.

However, if the banks are no longer paying the broker, somebody needs to pay the broker, which means that the consumers are going to have to pay a higher fee. Is this good or bad? Well, it depends, right? If I pay a higher fee but I get a much better product, I might be willing to do that as a consumer. However, if the previous allocation of mortgage was pretty good, then -- and I have to pay more for it, then it's not -- no longer a big deal.

There's also a tradeoff when it comes to the broker. Before what happened was that the challenger banks, the new banks, were paying a lot of money to the broker, and since the broker had higher costs for this bank, he was willing to still originate mortgage with the new banks because he was getting compensated for it. Now, the new banks don't have this tool to, in a way, bribe the broker or compensate them for higher cost. Therefore, what I see is that in the network, all the weight goes to the Big Six.

So now what you have is the following: Brokers

only do business with the Big Six. Banks going direct -- sorry, borrowers going direct, more often than not, go with the Big Six. So at the end of the day -- let me just show you -- yeah.

So at the end of the day, because of this multifalling competition, I find that interest rates go up by 11 percent. So since I'm out of time, let me just conclude. So essentially I think that in markets where there's high search costs, when consumers can go directly to the provider and the provider might exert some market power through the direct channel, banning payments to intermediaries may not be a good idea. Of course, this is UK market-specific, but I think the point of the paper is that we need to understand two things.

One is that there's always tradeoffs when you put in a regulation. There's positive side, there's negative side, and especially sometimes in a lot of previous work we focus on demand side reactions, but also supply side is going to react, and in the case of this market, when there's a network, the network is going to change, and that can have dramatic effects on competition. With that, I'll finish. Thank you very much.

(Applause.)

MR. VIOLETTE: Now we have Jean-Francois Houde from the University of Wisconsin to discuss.

MR. HOUDE: Okay. Thank you very much for having me. This is a great paper, I must say, and like, you know, many successful -- this was Claudia's (indiscernible) paper last year, if you know, and like many papers that are successful, it's really two or three papers. There's just a lot of stuff in that paper. So there's different ways of reading this paper. You could think of it as a finance paper, and that's a lot of -- you know, how Claudia's pitched it, whether there's agency problems and problems with intermediaries, more of an IO way of looking at this, which is a paper about vertical integration in some sense, right, the brokers are breaking the link, so that's kind of how I'm going to think about this.

So I'm -- you know, again, different ways of reading this paper. That's -- that's what I like about the paper. It's -- it really analyze the effect of competition between firms that have different degrees of vertical integration, okay? You can think of the brokers as the retail channel for banks, and banks are supplying to both channels, and we have this vertical structure.

And the goal of the paper in that sense, if you

think about removing brokers or regulating commissions, is basically trying to ask the question, well, does vertical integration have any kind of anticompetitive effect in this market and is banning wholesale price discrimination, which is banning, you know, dispersion of commission, good or bad for consumers, okay?

So, now, the data is amazing. The model is very nice. I'll talk a lot about the model, but the data is quite impressive. So I've worked on mortgages, and I've worked a little bit on brokers, but it's really hard to get data on the transaction price between those different areas and also just the fees that these brokers charge, so this data contains everything.

And so you see these upstream prices, you see the downstream prices, and as Claudia was saying, this is a simpler market to study than the U.S. market because, well, there's a single price. There's essentially no price discrimination across consumers, okay?

And there's also this -- again, as a vertical problem, there's also essentially a resale price maintenance. The banks are charging the same rate, whether you go to a broker or you go direct, and you also see the vertical network, which is -- so you see

for each broker what lenders they are dealing with.

Now, the model has a bunch of pieces, so, again, essentially there's no resale price maintenance, no price differences across consumers, but there's wholesale price discrimination, if you will. So the banks are going to charge different commission to different brokers, okay? There's -- you know, as Claudia mentioned, there's agency problems, and then we're going to relax the price-taking assumptions. So if you read the vertical integration literature, a lot of the papers are thinking about firms making different offers, and that's going to be relaxed here.

Okay. So this is how the market looks, okay? This might look different from the U.S., but it's actually not that different from what the U.S. market look like. So I've, you know, I've taken one of these challenger banks, pretty much random, might not be a good choice, but whatever. So you basically have HSBC, right, which is a vertically integrated company that sells both direct, and even these smaller institutions, right, who -- and in my picture, I've assumed that they don't have access to consumers, which is almost true for many of them, okay?

So essentially without brokers in this kind of market, HSBC would have a lot of market power because

consumers don't have access to these small institutions, and there's a lot of these small institutions, okay?

Again, this is -- this looks different from the U.S., but in the U.S., if you look at the vertical chain of mortgages, it looks a lot like this. We have a lot of -- not necessarily -- we don't necessarily call them brokers, but there's a lot of financial mortgage specialists in the U.S. that are -- and they originate more than half of the mortgages in the U.S., and they're getting their loans from somewhere, and they're not getting their loans from their own deposits, because they're shadow banks -- this is the term -- and so that's exactly how the U.S. looks like. It's just slightly different -- I mean, not exactly, but there's a lot of that going on.

Now, the model has a bunch of pieces, so, again, the rate is going to be the same, so that's very different from the U.S., so the rate is going to be the same by one or two. The consumers are going to pay a fee, and then these commissions, all right, are going to be allowed to be different, okay?

Now, what do brokers do? So the key thing -- and that's what's nice about the paper -- is that the brokers have a role here of allowing consumers to have

competition between the two upstream banks, okay? So without brokers you wouldn't have as much competition between the upstream banks, and that's what brokers do. They also reduce transaction costs, so there's kappa here, you pay a fee, so that kind of upset each other, but in general, at least in the model, you paid a higher fee -- higher transaction cost.

There's also efficiencies, so that's -- we're going to talk about that a little bit later. That's a little bit more ambiguous, but it does seem like, on average, the marginal cost from the bank is lower for the broker channel, and then there is this agency problem. So where does the agency come in? So agency is that I'm not going to solely give you the lowest rate available. I'm going to account for the commission, so in this model, I ignored a lot of the logit stuff, but that's basically what the role of data is.

Now, you can think of it as an agency cost, but another way of thinking about this is that theta plays a role of allowing this fringe bank to steal business, so by discount -- by giving a high commission, I can essentially steal business toward my -- if I don't have any branches, I can steal business. That's what data does, okay? It's basically the steering. You can call

it an agency cost, but that's basically what competition -- that's where competition enters.

Okay. Now, one thing, if you think about this as a vertical problem, there is no double markup here. There is no double marginalization. There kind of is, because the fee is chosen by the broker, so in principle, when I change the commission, the fee could adjust. That's assumed away here. It seems like the fee is zero for most, but that's a little bit of the weak point of the paper here.

Now, bottom line, vertical integration is bad, okay? So if you think about the equilibrium effect and overall, this competition channel is really important, okay, and that's the -- again, I'm oversimplifying things here, but if you think about the ways that people are paying, essentially this competition channel that (indiscernible) is really important. The same thing if you -- if you impose uniform commission. Price discrimination is good here because it allows this steering which then allows the small banks to compete more effectively with the big banks, okay? So it has this surprising upstream price discrimination as a pro-competitive effect.

Okay, so I'm going to be going through the blocks of the model fairly quickly. So this is about

mortgages, but really it looks like a model for gas stations, okay? So commission, I'm going direct, I'm basically drawing a bunch of logit charts, thinking about where to go, thinking about the branches, thinking about the characteristics of those products.

If I go indirect, I do the same thing, except that I'm steered, right? I have this bias that's coming from -- from the commission, and the deltas here of kind of the quantity of these products are viewed a little bit differently between brokers and nonbrokers, which means that in the model, the -- I'm not going to buy the same product if I go see a broker okay?

Now, the key assumption is that the search cost is the same, so it's going to cancel out, so it's not going to affect which bank I choose. The same thing with the fee, and like I said, that -- one question, there is no outside option here, so there must be a normalization somewhere. Anyway, I'm a little bit confused.

One implication -- I should have said more of an assumption -- here, there is no selection unobservables, so when I choose broker versus dealers, I don't think about this as a function of my taste for having a high LTV mortgage or a low LTV mortgage, right, which you might think is important. If I look

at my data set, brokers tend to deal with bad consumers, high FICO score, big loans.

Now, in the model, that's going to be -- that's going to be the bias, okay? So basically why is it that people take high LTV loans with brokers? It's because of this bias. It's because of -- now, the way I would interpret this is, well, if I go see a broker, maybe it's because I'm a little bit worried about qualifying for many banks, so I am going to use a broker to qualify for more banks. So bad consumers will select into brokers. So it's not clear to me how do we interpret this bias, whether it's selection or whether it's really a bias.

The other assumption is that, again, this is like a model for gas stations. So J here is a product, so it's a lender and an LTV. This is not the most attractive model if you think about demand for mortgages, because basically we're going to have I substitution patterns against different loan sites. If I think that the price of one loan is going to go up, maybe I'm going to substitute to buy a slightly bigger loan or slightly smaller loan. That's not really what's happening here because of the logit, so -- but I understand, you know, the reason, but...

Now, the price competition -- so price

competition looks pretty standard, so it's like a Bertrand model. The only wrinkle here is that there is these two efficiencies. One important part of the model is try to get the two efficiencies, the marginal cost of going to broker, marginal cost of going to a bank, now -- and this is going to sound a little bit like the comment earlier about the two -- the two of the unknowns. So here that means that I have J first for the conditions, and I have two times J unknown, so it's not -- it doesn't look feasible.

Now, Claudia does a nice trick here that she's not really going to estimate the true marginal cost. What she does is actually estimate the average marginal cost, okay, and then test whether the slope of these axes differ as the share of broker transactions, which in a bank changes. Now, those shares are not quiet exogenous, while they're depending on the price and they're depending on everything, so there's a little bit of an identification problem of how do you actually separate these two things, okay?

There's also the thing of -- you know, there's unobserved constructs, so the axes won't explain a lot of the variants, so where -- how do you assign the residual? Do you stick to -- and that matters for a lot of the counterfactual.

Okay, bargaining is -- again, it looks very standard. There's a wrinkle that she thinks about about participation, which a lot of previous literature doesn't really think too much about.

Now you are going to kick me out? Okay.

Now, there's one thing about the -- let me just mention this. So there's -- again, we are going to do the same kind of trick. We are going to invert a bunch of the conditions. The residual here is going to be the bargaining parameter, so this is a bit like Matt Grennan's earlier work. There's one thing -- the model is very rich, because here, this is a sequential move. So when I change the commission, I have a little bit of a raising rivals' cost effect or I want to soften price competition with my commission, and so there's a pass-through of the commission on the rate. So the first order condition is a little bit more complicated than what it sounds, and I'm not quite sure how it's handled in the papers. So there's probably more detail than this.

And it's also that it seems like participation is an issue, so I don't know if it's an issue at the estimated parameters. If it does, then it's actually that the Nash in Nash is actually consistent. If links are broken, then the others will react to this. The

Nash condition assumes that, you know, everybody has correct beliefs. So I have a bunch more comments, but then I'm going to be kicked out, so I'll let you go. Thank you.

(Applause.)

MS. ROBLES-GARCIA: So, thank you, Jean-Francois, like we can talk a lot more, like great points, and I was super happy when I knew you were my discussant, because -- yeah, we can talk more, yes.

Yes, go for it. I think people know there are drinks afterwards. There's a question there.

AUDIENCE: (Off mic) -- this issue. I think, bottom line, is that if you rolled the up-front costs in and paid for them with a yield spread premium and, therefore, you shopped only on the interest rate and the -- you know, how long the terms were for the loan, people did better. Are you able to examine this type of simplified search strategy using your data?

MS. ROBLES-GARCIA: So the good thing about my (indiscernible), which is different in the U.S. as opposed to the UK, is that people going to the broker, people going directly to the branch will get exactly the same interest rate. So the brokers do not get compensated based on the interest rate they get on the loan, while in the U.S., the broker commission is also

a fraction of the interest rate, while in my (indiscernible), it's a fraction of the loan amount.

AUDIENCE: (Off mic).

MS. ROBLES-GARCIA: So they say interest rates and fees, but usually the fees are zero or a thousand, and they can always be rolled into the loan. So they don't often pay the fee. So in my model, I collapse those things into, like, a net present value interest rate. One of my co-authors, he has a paper showing interest rates and fees, and the demand estimates do not change much, but -- yeah.

Yes?

AUDIENCE: It seems like the products are all pretty simple. Why isn't there a public agency that just lists the six banks, plus whatever entrants there are, with the -- there's only six interest rates, right?

MS. ROBLES-GARCIA: Yes. So on average, there are about 15 products per bank, because there's different loan-to-value bands, and then there's initial fixed periods, but I completely agree. Like, we have price comparison web pages, and why don't people search themselves? So you would imagine that the fact that there's these price comparison web pages, that they should bring the search costs down.

I think -- I couldn't go into the details, but what I call search costs have two components. One is the information gathering, so how costly it is for me to find the best product, and there's also the time component, how long it's going to take me to get the mortgage.

And what I see is that the very low-income people go to the broker, so you could imagine for these people it might be harder to search and to understand the products, but I also see the very high-income people going to the broker, which could go to the story of time. But, yeah, I completely take your point, yes.

AUDIENCE: We need a common application.

MS. ROBLES-GARCIA: Yeah, exactly. Okay, I think I'm done. Yes.

(Applause.)

MS. DUTTA: So, thank you, Claudia, and everybody else who participated in this session. We are done for the day. There are drinks outside and some food. We do have to leave this particular conference room by 6:00 p.m., so please make sure to take your stuff outside, but you can stick around in the reception area for longer than that. So, please join us at the reception.

Tomorrow morning, we will be starting at 8:45.

So we hope to see you all there.

(Whereupon, at 12:25 p.m., the conference was adjourned.)

CERTIFICATE OF REPORTER

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

s/Linda Metcalf
LINDA METCALF, CER