FEDERAL TRADE COMMISSION

I N D E X

FEDERAL TRADE COMMISSION

FEDERAL TRADE COMMISSION)
MICROECONOMICS CONFERENCE) Matter No. P085800
---)

Thursday, November 3, 2011

FTC Conference Center
601 New Jersey Avenue, NW
Washington, DC 20001

The above-entitled hearing was held, pursuant to notice, at 9:00 a.m.

P R O C E E D I N G S

- - - - -

DR. CHESNES: Good morning, everyone. Let me be the first one to welcome you to The Fourth Annual Microeconomics Conference here at the FTC. Thank you all for coming.

I'm Matthew Chesnes. I'm an economist here at the FTC. And before we get started, I just want to mention a few logistical items.

The conference is being recorded by a stenographer in the back. So, for the presenters up here, just try to speak into the microphone, and during Q&As, we'll have some roaming microphones throughout the audiences. Try to use those when you're asking questions.

There's evaluation forms which you should have picked up on the way in. So, try to turn those in by the end of the conference, just kind of -- it's a good way for us to improve the conference in future years.

The restrooms are located out by the lobby. If you go out to the left of the guard's desk, you will see a sign that will point you in the right direction.

There is Internet available in this conference room. There is a pamphlet available at the front desk, if you didn't pick it up, that has the code to get

online.

And then, finally, a security briefing. Just keep in mind, if you go outside the doors, you will have to go back through security on the way back in, which can take a few minutes. So, if you go outside during breaks, just keep that in mind, that it's going to take a few minutes to get back inside.

In the event of a fire or evacuation or earthquake, just, you know, exit in an orderly fashion. Our rallying point is across the street from the FTC. The Georgetown Law Center is right across the street. So, that's where we'll -- if you cross New Jersey Avenue, there will be someone there that you can check in. In the event that it's safe to remain inside, we will give instructions as to where to go. And then, finally, if you spot any suspicious activity, please alert security. I don't know what that means.

All right. So, it's now my pleasure to introduce the FTC's very own Joseph Farrell. Joe received his BA and doctorate from Oxford University and went on to teach at the University of California at San Diego, at MIT, before joining the faculty at Berkeley in 1986. At Berkeley, he was a professor of economics, Chair of the Competition Policy Center, and affiliate and professor in the Haas School of Business.

Joe has also served as chief economist for both the Antitrust Division of the DOJ and the FCC, and Joe has been Director of the Bureau of Economics here at the FTC since 2009. So, please join me in welcoming Joe Farrell.

WELCOME AND OPENING REMARKS

DR. FARRELL: Thank you, Matthew.

This is a slightly unusual agency in that while we do policy, it's mostly a law enforcement agency, but the laws that we enforce mostly say do good economics. So, the relationship between the day-to-day job and thinking about economics is closer than it is in a lot of law enforcement.

And one of the things, as Matt mentioned, that I've tried to do in my work is to go back and forth between the academic and policy worlds, and one of the things that we in the Bureau try to do is to keep in touch with the academic world, to leverage the academic world.

I edited an economics journal for five years and was much struck by how many very smart, hard-working academics there are who have no idea what an interesting problem is, and it seems like the co-existence of that problem with the fact that there are a lot of interesting and important problems but not enough people

working on them is something that it ought to be possible to do something about. And so, through a variety of means, we try to do something about those things.

One of the things we do is we have a research program here in the Bureau of Economics. It's pretty active considering that everyone has a day job. In the last year, since the third annual conference, over 20 papers have been completed and released by BE staff, and over 20 have been published -- not, of course, the same 20, and I'm talking to you, journal editors -- but some of the publications have been in the American Economic Review, the Journal of Industrial Economics, and we apparently completely took over the International Journal of the Economics of Business for a while. So, that's one of the things we do, and we've had some important successes there, and it's a great part of our mission.

Obviously another important part is this conference, where we try to bring together top academics with our own staff and share our research, try to convey a sense of what we're working on, what we wish you would be working on a bit more, and other such varied agenda items. Obviously, the bulk of the discussion is going to be about actually the specific problems that are

addressed in the papers, and I'm looking forward to it.

I'd like to take this opportunity to thank the scientific committee, Mark Armstrong, David Dranove, Aviv Nevo, and Nancy Rose, and some of the staff who helped put this conference together, which is a great deal of work. Laura Kmitch, who is responsible for most of the logistical arrangements, so if anything works, thank Laura; and the economists Chris Adams, Matt Chesnes, who you just heard from, Patrick McAlvanah, Jason O'Connor, and Dan Becker. Finally, I would like to thank the Searle Center for sponsoring the lunch and reception at the end of day.

And so although I haven't used, I think, all the time available, I'm sure we'll find a way to make up for it. So, let's get moving.

(Applause.)

DR. CHESNES: Thanks, Joe.

This morning, I have the pleasure to introduce to you our first keynote address speaker, Aviv Nevo from Northwestern University. Aviv spent his time at Berkeley and MIT before moving to Northwestern in 2004. He is currently a professor in the Department of Economics and the Marketing Department at the Kellogg School of Business, as well as a research associate with the National Bureau of Economic Research.

He received his BA from Tel Aviv University and his Ph.D. from Harvard. I would like to thank Aviv for being a part of this conference for the last four years and for being on the scientific committee for the last three. Please join me in welcoming Aviv Nevo.

(Applause.)

KEYNOTE ADDRESS

DR. NEVO: So, it's a bit intimidating now to talk, because first Matthew tells us that any suspicious activity should be reported to -- was it the authorities or security? So, I'm afraid that, you know, when I get to Table 1, you know, someone is going to run out and report, you know, "There's endogeneity in the table," and the next thing I know, I'm in detention.

And then Joe tells us we have a bunch of interesting economists or smart economists working on uninteresting problems, and I'm, like, if I had the guts, at the end of the talk, I'd ask, "Well, did I pass the test?" But I don't have the guts to ask, so I'm not going to ask you that, Joe, at least not in public.

Anyway, so I thought today -- kind of over the past few years, I've given different versions of what we'd call a keynote address. So, here I'm going to talk about, you know, some research I've been interested in lately, and it's going to be basically based on mostly one paper, joint with Rachel Griffith and Pierre Dubois, but actually on kind of a broader research agenda. And so it's going to be somewhere between, you know, discussing on kind of a general topic and one specific paper.

So, what I'm going to focus on is -- you can see

the title here is, you know, "We Are What We Eat," but "What Do We Eat and Why?" So, let me start with a little bit of motivation.

Obesity is linked to a bunch of bad things, okay? I could go through this whole list, but the bottom line is obesity is bad, at least for our kind of personal health. There is estimates of the costs of obesity, and they range from, you know, $147 billion to $168 billion in the U.S. alone, much more globally.

One thing that I find actually quite striking about obesity -- and actually, when -- let me just take a second to stop here. When I talk about obesity, so we all -- at least when I started this project, I had this image of obesity. I had this image in my mind. I drove across country about, you know, 13 years ago, and I think we stopped at an Indiana State Fair, and I think I have somewhere a photo of three relatively young ladies, I would say kind of probably in their twenties, kind of walking toward the fair. So, it was kind of taken from behind them. And, you know, I had to use a wide lens to fit them all, I mean, so that's what we have in mind, right, these people that we hate whenever they kind of sit next to us.

But the truth is, that's not what obesity is or the definition of obesity. Obesity is defined really by

your BMI, which is basically your weight divided by your height, and obese is anyone who has a BMI of greater than 30. And I don't know if you guys have ever looked up your BMI -- if not, I'm sure you will after this talk -- but, you know, it's not that hard to get over 30. I mean, I like to joke and say that I'm, you know, one good Thanksgiving dinner away from obesity, and that's not really that far from true. So, you could model and you could make a decision to yourself that that's what you want, but basically, that's just what we use here.

So, obesity rates vary widely across many dimensions. So, first they're going to vary widely across different countries. So, for example, in France, the obesity rate is 14 percent, and my co-author, Pierre Dubois, tells me that, you know, they're outraged because it crossed 10 percent a few years ago, although I tell him, you know, if in this country we were able to get it down to anywhere near the French levels, it's probably going to be the biggest public health success any time, anywhere.

The U.K. is somewhere in between, and the U.S., you know, kind of -- the exact number varies a bit, depends how you measure it, but somewhere around 30 percent. I think, you know, Mexico is also kind of constantly competing as to who has the highest obesity

rate.

It also varies a lot across states. So, first order, it's basically the South and the Midwest versus the coast, and you can see some numbers here. So, Mississippi and Alabama, it's well over 30 percent, while in Connecticut or California, it's kind of in the low twenties, okay?

It also varies across demographic groups. Obviously, it varies by age, I mean, that's just a natural process, but it varies a lot by income and by race, okay? I don't have the numbers here, but that is. . .

Now, a key cause of obesity -- I mean, there's a lot of things that go in, but it's basically -- the first order, it's a very simple formula. You know, calories come in, calories that go out, right? So, if you want to reduce your weight, you either have to eat less or expend more calories, exercise more, be more active. That's the first order. That's really what it's all about.

So, we're going to actually be looking at the first side of this, kind of what's going to go in, what are people sort of consuming. So, in terms of general questions -- we're not going to sort of approach it, but there's a lot of general questions that come to mind

here from an economist's point of view. You could ask, you know, what's going wrong? And actually, even by saying that it's wrong, we have already imposed some sort of norm here.

Is it -- you know, are there kind of issues with prices? Are the relative prices wrong? Are the product offerings -- do people just not have, you know, the right products or the good -- the healthy products available to them? Is it a matter of preferences?

I mean, maybe, you know, we say -- and I'm sure we've all been in that situation, you know, you have this kind of heavenly dessert, and as you're eating it, you know your life expectancy is going down by a week, but you're saying, "But it's worth it," okay? I've been in that situation. I'm sure you have, too, right? So, it's a rational decision, say, "You know what? All right, I'll live a week less, but, man, am I going to enjoy it." Or is it some sort of behavioral story, right, that there's some sort of dynamic inconsistencies, right? There's all these kind of stories going in.

You could ask, you know, questions, is there any role for government intervention, right? Is there externalities or is this sort of a private problem, right? If people want to, you know, kill themselves by

overeating, why should the government care? What are the externalities here? To the extent that government does care, what are the type of policies that might be effective?

I actually find that, you know, in this realm, it seems like -- you know, maybe it's because we're not quite sure the government should really be in this business, but it seems the policies that people are pursuing are not that innovative. It's really more about, well, that's just tax stuff, right?

You think, for example, you know, when it comes to environmental stuff, we think, you know, something like the CAFE standards, right? We're going to force automobile manufacturers to have certain standards for the cars they sell.

Suppose you did that for supermarkets? You said, you know, for every -- every two-liter bottle of Coke you sell, you also have to sell a liter of broccoli -- or liter -- a pound of broccoli, okay, or you have to sell -- you know, your overall content of what you sell, your nutritional content has to satisfy the following conditions, okay? Why not, you know, do that?

Now, maybe it's because we don't really think that government should be in this business, or some

limiting, but to the extent that we do -- and it seems like some people do, right? I mean, the White House or at least our First Lady thinks we should be in this business. I think there's a lot more that we should look at.

Now, of course, I'm not going to be talking about any of these, but this is just kind of general questions. So, what am I going to talk about today? I'm only going to focus on the across-country difference, I think with the exception of one slide, and that's all that I'm going to show here, and my goals are the following:

First, I'm going to just provide you some descriptive statistics, right? So, when -- you know, if you think of -- the first question, what is it that we eat? Well, let me show you what the difference is across countries, right?

So, we all have the image of what it is that the French eat, what is it that the U.S. consumers eat. Well, let me actually show you from real data. So, that, I think, is the first part and maybe kind of the major part of what I'm going to talk about.

Then, what we're going to try here is a little bit more ambitious, which what we're going to try to do is we're going to try and separate the differences --

the observed differences in the consumptions into differences that are due to -- I'm going to call it something vague -- the economic environment, okay? And I think what I have in mind here are things like prices and product offerings. So, I'm going to try and separate that versus what I'm going to call other factors. You can think of it as broadly defined preferences, okay?

And what we're going to do is we're going to have very detailed data at the household level of basically everything that was purchased over a fairly long time period, right? I mean, at least a year; in some cases, we followed these households for two or three years. So, basically, you can think of we have everything that they purchased in a supermarket, okay? We're going to match that with very detailed product-level information, okay? So, basically, at the bar code level, we know everything that appears on the label, okay?

So, we can actually tell you this is how many calories you brought into the home. This is how many grams of protein or carbs or fat you brought into the home over the year. I'm going to treat it as if you consume, okay, although one of the things you could say, "Well, what happens if we buy all this stuff and we're

throwing it away," okay? There's a lot of caveat that I'm not sure we'll get to, but -- so, we're going to have that data, and we're going to have it for three different countries, for France, the U.K., and the U.S. So, I'm going to provide the descriptive.

Then we're going to use that to estimate a demand system at a very kind of aggregate level. So, it's not going to be demand for a two-liter bottle of Coke; it's going to be a demand for food. And I'll show you how we plan to model that.

And then we're going to try to simulate some counterfactual consumption. So, the thought experiment that we're going to try to do is to say, "Well, what happens if you're going to take a U.S. consumer and put them in France?"

Now, it's not literally put them in France. It's basically, say, take a U.S. consumer and give them the French "economic environment," so French prices and French food offerings, okay? That's kind of the thought of -- you know, if you want, we're either bringing France to the U.S. or taking a U.S. consumer and putting them in France, okay? So, choose whichever one you want to think about. That's what we're going to try and simulate.

Okay. So, let me just tell you a little bit

about the data, because it's really all driven by that. So, these are data collected by marketing companies. In the U.S., it's Nielsen, the same Nielsen from the data rating, from scanner data, we've seen a lot of this, but here, it's collected different. It's not collected at the store. It's collected at the home. So, it's called -- it's a home scan data.

So, what they do is they take a panel of consumers, they give them scanning devices that they scan, and they're supposed to basically, after they buy this and come home, just like -- you know, just like you might at the self check-out counter in supermarkets scan everything, they do that all here. All right, for things that doesn't have a bar code, they have a code that they're supposed to punch in. And these households are actually pretty quick at doing these, and those that sign actually tend to like doing this stuff.

They record everything that's purchased. We know the exact date and location of purchase. So, we know exactly they went to, you know, a particular supermarket on exactly what day and the location. We know the quantity and the price. There's a bit of an issue how the price is computed. For the most part, it's actually -- what Nielsen does is after they find the location and the item, they actually go to the store

to get the data. So, the consumers don't actually have to punch it in themselves. That's -- at least in the first order.

In total -- it depends on how you define observation here, right, but if you define observation as an item on a particular purchase, you know, we have literally hundreds of millions of transactions. We have very detailed demographic information, which varies a little bit across the country, and then we also have nutritional information -- which in the U.K. comes with the data. So, there it's not -- it's not Nielsen, it's a company called Kantar, but they actually collect it.

In France, one of my co-authors collected this by hand. I mean, he basically had a team of RAs that went and documented everything from supermarkets. And as a result, we only have a small number of things, just what we call the macronutrients, which is what I'm going to focus on here. And in the U.S., it actually comes from a company called Gladson that was actually purchased by the USDA, and it's very, very detailed.

I mean, if you look, it has something like, you know, a few thousand different attributes, although most of them don't appear, right? So, it's literally everything that's on the label, right, including all the little fine print, right, that we never bother to look

at. Okay. So, it's very, very detailed.

You know, I'm not going to bore you with all the details of how we did the matching. You know, even just doing the across-country matching, you might think it's -- well, it's trivial, right? Food is food. It ends up it's not. It's quite amazing, and that's kind of part of, I think, the big thing here, right? If you look at the different categories that you have in, you know, the U.S. versus the U.K. even, right?

I mean, France, you might say, "Oh, you know, the French, of course, they're different, you know, but U.S. and U.K., you would think they are really similar." It ends up not. Just the fruit categories -- you know, a big item in the U.K. is baked beans, okay? It ends up in the U.S., we can barely find baked beans. They eat a little bit of it in Boston, but that's about it. We eat a lot of refried beans, which they don't have in the U.K. at all, right? But -- so, just doing this matching, right?

So, you know, it's -- the categories are different. Within categories -- if you look at yogurt, you know, yogurt, that's a pretty standardized thing. No, right? I mean, the types of yogurts that we have here, even though it's the same companies, they offer completely different products here than they do in the

U.K., not to mention France.

So, anyway, we had to do a lot of data work here. Basically, what we did is we used some classifications used by the USDA that has 52 product categories, and we actually collapsed it basically to nine different broad categories, which I'll sort of show you today, and even at this level, there was actually sometimes issues of where we put what.

So, let me show you just a little bit about the sample. So, this is the sample we have across the countries. The U.S. is actually the smallest sample for reasons I didn't talk about. It's because we want to include the purchase of fresh food.

One thing that's noticeable about the sample, it tends to be relatively older compared to the population. So, there is a little bit of sample selection of who participates, right? So, we all think in our mind, you know, who in their right mind would -- we're thinking about ourselves. Well, it's true that we're -- we or people like us are not in the sample, but it's also true that we're irrelevant, okay, in the sense that high-income people are really not in the sample, but that's -- you know, we're -- you know, if you look at the whole, we're not kind of important.

What's really undersampled here, actually, lower

income single moms, black families. Oversampled, white, older, single, but, like, widowed women would actually be probably oversampled in this, and you can see this when you look at the descriptive statistics. The overall age is high, the number of kids is lower, and the household size is smaller. But that's if you want the bad news.

The good news is it actually seems to be roughly comparable across the countries, and we could try to weigh these to try to make them more representative. We're not going to do any of that. This is just the raw data. It is what it is. So, a comparison between these samples. We don't want to fudge with it.

Okay. So, let me kind of dive in and just give you the numbers. So, this is just kind of at the aggregate, across the countries. So, what we have here is these are daily -- let me just first, on that first row, daily calories per person, or what we have is what we call adult equivalent, so we kind of -- to basically control for the different composition of the household, there is a required or kind of suggested number of calories that are a function of your gender and your age, right?

So, to do that, what we do is for each household we compute what is called a household -- an adult

equivalent, right? So, we kind of translate into what kind of an adult would consume, and then we -- these are an adult equivalent. So, the recommended amount for the adult equivalent would be 2500 calories per day, and you look at these.

So, first, all the -- low, but remember, these are all food purchased at home. It does not include food purchased outside the home, and it does not include alcohol, okay? So, you know, don't be kind of excited about, "Oh, look, we're great. We're -- you know, we're below the level. We're doing great." It's not quite that.

The important thing is to actually here compare across countries and maybe kind of matching our prior -- you know, you'll see France has the lowest number of calories, the U.K. in the middle, and the U.S. higher. And, you know, these differences are quite significant, right? So, 100 calories per day, if you translate that, that's -- you know, that's about ten pounds a year, okay? That's a lot.

Now, maybe more interesting than the total amount is where these come from. So, if you look at the next three rows, we kind of broke them up into calories that come from carbs, from protein, and from fat, and the percentage that come. So, what's interesting here

is that the U.S. and the U.K. are almost identical, and interestingly, they're actually pretty close to the recommended amount.

Again, recommended amounts for overall consumption, this is just for food at home, right? So, to the extent you think food outside the home tends to be with more fat, then maybe, you know, we're exceeding, but at least in the home, we're actually pretty close.

The French are actually a bit different, and here maybe it's -- well, maybe after you think of it, it kind of makes sense, but if you look, their percent fat -- even though they're overall much lower in calories, they get a lot of their calories from fat, right? It's all that cheese. And, actually, you'll see in a second here, I'm not joking. It really is all that cheese.

On the other hand, where the U.S. really gets clobbered is on the carbs, okay? You look at sort of the amount, and if you look at the next row, just the amount of carbs, right, the -- in grams, and that's, again, something that we'll see over and over and over again. That's where we're getting clobbered, right? And you'll see where we tend to consume categories that have more carbs, like soft drinks, and then within these categories, products -- you know, we just have to have

products that tend to have more carbs in them.

This is the one thing where I'm going to deviate, actually, from looking across countries. So, I just have -- here's an equivalent table, a similar table, just looking in the U.S. across cities. I'm just going to show this once, but just to kind of motivate other work.

So, if you look across cities, so I just gave you represents, so I have here Chicago, LA, San Francisco, Philadelphia, and then Midwest is a collection of cities. The sample is too small to use any particular Midwest city, so I kind of collected it. So, these are places like Columbus, Indianapolis, Grand Rapids, I think Minneapolis is in here, and the same for southern cities.

And you can see what we -- you know, kind of what we saw along the obesity rate, but it's reflected here as well. LA and San Francisco, in terms of their overall, they're much closer to the French, right, just in terms of calories. I didn't give you the breakdown. Chicago is somewhere at the national average, and then the other kind of smaller -- I mean, these are not rural cities, the smaller Midwest cities are much higher, as are the southern cities. So, you see this kind of huge regional variation.

And, once again, you kind of ask, you know, is this driven by preferences or is this driven by the environment? And, actually, in the U.S., we're able to do something a little bit more than what we can do across countries, because in the U.S. -- and I am not going to report on this today, because we just got the data -- we actually have migration data.

So, we actually get to see someone that lived in Columbus and moved to San Francisco. Now, you might say, "Well, it's not random that this person moved," okay, but we get to see in some cases literally their consumption in Columbus and then in San Francisco, but also, we get to see kind of, you know, how different they are relative to people in San Francisco. And I don't have those results yet to report on them, but that is something that we're looking at in the future.

Okay. So, that was just kind of a detour. Let me go back to the across-country comparison. So, here, again, comparing across the three countries by the nine different categories. So, what I have here is the first three columns are expenditure, the next are expenditure shares, and then it's quantities. And here, by quantities, we have to define them somehow, so we define them in kilos, okay?

Okay. So, basically it's kilos of fruit or

kilos of meat. Ideally, I would have liked to define this as servings, right, to the extent that that's maybe capturing -- in some categories, we will say, "Well, kilos are a good representation, in some not," but we can't. I mean, we can do it in the U.S., but in France and the U.K., we just don't have servings. In the U.S., we have it off the label, the recommended amount. We can't do that. So, we have to go and resort to doing it in kilos.

So, what do we have here? So, let me actually focus on the expenditure shares. What's quite interesting here is that the expenditure shares -- well, first, if you look, the U.S. and the U.K., just like we saw overall in the aggregate, tend to be much more similar. France is really the outlier in these. So, that's kind of the general theme of what we find.

So, if we go down, for example -- let me just highlight a few of these numbers. If you go -- if you look at meat, for example, the French spend 31 percent of their expenditure on meat. They also spend a lot on dairy, 17 percent. On the other hand, they spend much, much less on prepared food, okay? These are TV dinners or frozen pizzas or canned soups, okay, while the U.S. and the U.K. are actually quite similar in these. Now, of course, the U.S. and the U.K. are not identical, but

they're actually much, much closer to each other.

Now, why am I focusing on expenditure shares? If you think, for example, of the Cobb-Douglas utility function, well, these are the utility parameters, right? I mean, these are basically -- you know, think of -- note them as alphas, right, how much you weigh on each of these products. These are -- so, in that very limited and restricted utility, that's the preferences.

Then you say, "Well, how about prices?" Well, let's go to quantities, right? So, now, when we look at quantities, here, what's interesting is here things are actually going to change a bit. So, now, if you look, for example, at meat, you know, there's a huge difference in the expenditure shares, but in quantity, actually, here, France and the U.S. is almost identical in terms of the kilos, how many kilos of meat we consume. This is, by the way, per quarter, okay, so, how much we consume.

So, now suddenly, you know, even though the U.S. and the U.K. might actually seem quite similar in terms of their preferences, if you believe the Cobb-Douglas function, right, their actual consumption looks a little bit different, right? So, that's, if you want, kind of -- if you want, like, the simple version of what we're going to do, is we're going to do something that's

only marginally fancier than this, because we are going to have -- actually, we are going to build on some kind of a Cobb-Douglas-like function, but, you know, change it a little bit. At least, you know, for now, that's basically the whole story right here, right?

You know, the middle set of columns, that's preferences. Once you interact, you get the quantities. Now, of course, the quantities now you have to interact with what are the nutritional content, and I'll show you those in a second.

So, let's go directly -- I mean, you could indirectly infer the prices from that comparison, but you could actually look at the prices here, and you can see there's actually a fair bit of difference in the level of prices. What I don't actually have here is the relative level of prices. So, you know, again, there's a lot of numbers here.

Let me just focus, for example, on drinks. So, if you look at drinks, France actually here is the cheapest, followed by the U.S., and then the U.K. one of the things -- let me go back for a second -- I didn't highlight here is if you look at the expenditure share on drinks, that's one of the places where the U.S. is clearly an outlier, okay? So, here, France and the U.K. are similar, but this is actually high -- this is

actually missing an important point here.

What the drinks are are different, very. So, in France, drinks basically mean or mostly means bottled water, okay? In the U.S., it's soft drinks, okay, and that's actually different. So, the price in the U.S. is actually lower than in France. Bottled water in France is cheaper by, you know, one or two cents, but -- not by much, but soft drinks are cheaper, all the other drinks are cheaper as well. It's just here it's a matter of, you know, you're aggregating different things, okay?

So, drinks in the U.S. are actually significantly cheaper than everywhere else, and, indeed, we actually consume more, okay? Now, what's kind of -- what's driving what, I don't know. There's a big endogeneity issue, so before you go reporting me to the authorities, I'll admit it right away, but -- but just kind of as a descriptive.

In terms of the nutritional content, okay, so these are basically nutritional by category, by country, and the general theme here -- again, I can go number by number. The general theme here is kind of what I said before, is the fact that if you look within each category, so not only do we consume more categories that are high in carbs, like drinks, within each category, we tend to have more carbs. So, if you look, for

example -- you know, you go down this list, compare the U.S. and U.K., you can see almost each and every one, we have more carbs in the U.S.

Now, the way we've computed here -- so, you could say, "Well, all this is saying is that" -- well, even within categories, right, you're just -- tend to go to the more higher -- higher carbs products. Well, the way we computed this here is actually it's not weighed by purchases. What we did was we gave one -- you know, an observation here is just a product, okay, not how many times you bought it.

So, if you think that there is Coke and Diet Coke, okay, it's not that you say, "Okay, we're just buying more of the Coke than the Diet Coke, that's why it's more in carbs." It's just that Coke is one observation or -- it's not really Coke. It's -- it's the UPC has one observation, and diet is, right? So, we're actually trying to take the weights out of it.

So, we're just trying to in some way give you the description of this is the universe of available products, okay? So, an observation here is a product that was bought at least once, and if it was bought once or bought 10,000 times, it gets a weight of one in this regression, okay?

So, to the extent that you say, "Okay, Coke gets

more weight," it's because Coke might have -- in the U.S. might have 30 different UPCs, and in France, only one, as opposed to Diet Coke, right? So, we still haven't kind of eliminated that margin.

But it is telling you that the product offering, okay, the universe of products has more carbs in the U.S. I'm just focusing on carbs because that's really kind of a big difference, but you could sort of see the same thing -- actually, let me just point out, you know, before I was talking about the French and fat, and you can see, for example, you know, you look at dairy, okay? The U.K. and France are similar, but you can see the dairy in France is a lot richer, a lot more fat than the dairy in the U.S., right, looking at the last three columns.

On the other hand, meats, it's the opposite. Meats in the U.S. tend to have more fat and actually less protein, so maybe are kind of lower quality, right? We like to eat our ribs. They like to eat their filet mignon, just kind of as a...

Okay. So, that was kind of the descriptive statistics. So, let me --

How am I doing for time? Five, okay. All right. Did I get any of Joe's extra time?

AUDIENCE SPEAKER: Yes.

DR. NEVO: Okay, good. Bargaining works.

So, how do we model this? So, the key challenge is we have this extremely rich data, and we have to decide how we're going to model it. So, there's two options. One is to estimate what I'm going to call demand the usual way, right? So, that's basically let's focus on a particular category, right -- cereal, yogurt, soft drinks -- and just focus that across the three different countries.

We didn't do that for a couple of reasons -- I can go through all of them, whatever -- I think in part because we were kind of bored with that, we wanted to do something different, but I don't know if that's the most, but maybe that's the truth.

So, instead, what we're going to do is we're going to go to a much more aggregate product definition. We're going to define a product as one of those nine categories. We could go to something a little bit more desegregated, but it's going to be at an aggregated level. And the key question here is, how do we model the -- how do we kind of bring in the nutrient information, right? How do we bring that into here, right? So, we know how to do this in kind of a discrete choice model. That's something we've been doing for a while. So, the question is how to do that.

So, we are going to offer here kind of a new demand system, okay, and by new, I mean we're going to go back about 55 years and use an idea that was actually proposed in 1956 by Gorman, actually in a discussion to a paper about estimating demand for eggs. This paper was actually published in 1980 in the Review of Economic Studies, although it was the actual original discussion by Gorman.

I highly recommend that you read the paper, one of these pieces of classic masterpieces. Whether you care about eggs or demand sets or not, I mean, this is just -- it's a beautiful paper.

So, we're going to kind of take that and bring a little bit of kind of modern notation and modern view and we're going to have a slightly different objective than the original paper, but a lot of our ideas are basically from that original paper. So, let me just kind of give you a highlight of what we're trying to do here.

Chris was telling me before you're not supposed to have equations in a keynote address, but report me to the authorities, once again.

Okay. So, there's end products. Each product is characterized by a C characteristic, okay? Think of these as the macronutrients. There's going to be a

numerator, xi. zis are the characteristics, and yis are the quantities consumed of these -- you know, think of these nine categories, okay? We're just going to put -- this matrix A is all the categories.

So, the consumer's problem is to maximize this utility function, and that's really kind of the new part of it. So, the utility function is going to have the numerator in it, it's going to have the y's in it, the products, and these are kind of the classic or neoclassical demand -- utility functions. But it's also going to have the z's in it. It's going to have both characteristics and quantities, all right?

Now, we're used to having z's in it, right, if you've followed any kind of discrete choice literature, and we're also going to put in, you know, the y's, and that's kind of the -- and you maximize that subject to a budget constraint and this production constraint.

Okay. So, you can just use standard first-order conditions to drive. So, basically what happens here is with this linear technology, if we didn't have the y's in here, you would at most consume C products, okay? So, once you think of it, because of the -- it's quite obvious, right? It basically says that -- think of if there's just one product. You'll say, "Well, what's the one that delivers, you know, calories in the most

cost-efficient way?" That's the product that I would consume, right? And that extends kind of more broadly. So, that's why we need the y's in here, and as a result, we can actually nest all the different products.

So, I'm just going to kind of rush through these. So, what we're going to have here is we're going to basically have a Cobb-Douglas in the nine categories, then plus some kind of function that also brings in the nutrients, okay? And I'm kind of going to skip some of these.

The great thing about this is that it gives a very simple estimating equation, which I have here in the middle of the slide. So, it's basically the total expenditure by individual, I, on this product in the -- in this period, which is going to be quarter, as a function of the characteristics, and then all this fixed-effects stuff, they are going to pick up -- on a bunch of stuff that are going to be fixed effects, they are going to pick up some variation in the preferences, but also some variation comes from unobserved product qualities, right? So, it's going to -- the regression is going to be an expenditures on the nutrient.

We have an endogeneity problem, because, you know, the more you consume of a product, your expenditure goes up, but so does, you know, the z's that

I have in here. So, we need an instrument, and what we're going to use -- again, I don't have time to talk about the details of it. What we're going to use is the availability of products.

So, we're going to assume that the availability of products, conditioned on all the fixed effects and everything that we have in there, is exogenous, right, so that's kind of an idea that's been used a lot in IO, and that's going to impact your choices, but in some sense on the -- you know, I always get this wrong -- the extensive/intensive margin, right? So, we're kind of taking that out, okay?

I'm not sure if any of that made sense, but I can clarify it later. These are the demand estimates. Let me just sort of point out one thing here. All IV -- let's look at the IV. These are the coefficients of how much you care about each of the nutrients. Let me just point out one number to give you an idea of why we're even doing this. If you look at the carbs, right, so the U.S. consumes more carbs but actually has a marginal -- a lower marginal utility from carbs, right?

So, to the extent that you believe these numbers -- and you might say, "Well, we don't believe the" -- but if you believe these numbers, this is sort of telling you that even though we're consuming more

carbs, it's not because we prefer them more. It's because of the environment, right? That's kind of what's available to us. And that's a very qualified sort of statement here. I don't want you to go and say, "This is a done deal," but that's just based on our numbers.

So, let me just kind of slip down to the counterfactuals, and what we're kind of simulating here is the effect of -- I call it an American in Paris. So, what happens if you take an American and put them in a different country? So, first you see -- let's just look -- let's just focus on France. You see their total consumption of calories goes up.

There's an issue here that I skipped over, which is the price of the outside good. So, we might need to do a little bit of rescaling, which you do in the last two columns, and once we do that, you'll see -- because initially, it actually seems like, you know, they're dropping their consumption by 400 calories.

That's huge, okay, which seems a little bit too much, and it ends up it's just rescaling of kind of, if you want, a total -- it's not quite an income effect, but the -- once we kind of do this -- and I didn't have time to talk about how we do that -- you see their consumption goes down, but actually now by about 85

calories per day, which actually makes sense, right?

So, it's saying, well, even with the preferences, right -- so, it's -- you know, so the environment is basically playing a key role. So, there is a role for preferences. So, they are still higher than what would be in this simulation the French numbers, okay, but the environment is playing a key role.

So, if you take a U.S. consumer, put them in France, their total consumption goes up but doesn't quite hit all the way down to what the French would consume. And we can look here and kind of see exactly where that's coming from, but I'm out of time. So, let me conclude.

So, we document a difference in food purchases, we estimate a demand model, used it to look at the behavior across countries. These very preliminary results suggest that economic factors are at play and are important.

What are we planning to do in the future? A, play a lot more with this, but also look at other dimensions, for example, the U.S. cross-regional differences, and here, what's really interesting is we do have the migration data, so we literally have actually people moving.

So, both we can kind of describe what happens in the data but also see things like, you know, do their preferences change as they live, right? So, as you move to San Francisco, do you initially purchase like you're living in Columbus, but then you drift closer away?

Kind of the basic idea is similar to a paper we actually had here, I think it was last year, that Matt Gentzkow presented where they looked at brand preferences, right, how much -- you know, what mayonnaise you buy, how that drifts.

Here we're going to do it but not at a brand level, but as an aggregate. It's actually -- the data we got is from them. And similarly, kind of for differences across demographic groups.

The data is not rich enough, but we could even sort of do things like sort of ask, you know, what would happen if you take someone from a food desert and put them in a place where there's more -- we don't have -- the data is not rich enough, so I am a little bit reluctant to say we can do that, but that's the kind of things that we could, in principle, try to look at.

Okay, I'm way over my time. So, let me stop here.

(Applause.)

DR. CHESNES: Do we have time for questions?

Yes, we have time for a quick question. Give Joe the mic.

AUDIENCE SPEAKER: So, you handle food eaten outside the home by assuming it's zero, essentially, or constant across the entire sample. Wouldn't it be a little better to take the national averages in the absence of any better data? And that would enable you to account for some of the differences in habits of eating outside the home.

DR. NEVO: Right. Now, let me stress one thing. At least across countries, at least as far as we can tell right now, the U.S., both in terms of expenditure and amount of calories eaten outside of the home, is higher. So, to that extent, right, to the extent, you know, if we want to compare -- you know, our current results, we're being "conservative," right, so the differences are even bigger.

One of the things we're actually planning to do is there's actually other surveys that are quite detailed about food eaten outside the home, and we were hoping to actually have a model a little bit richer than what you had, which is based on demographics, so based on your location and your demographic, to try to -- our main problem is we don't actually have prices, so, you know, we have to figure out, you know, how do we get,

you know, prices from McDonald's, but -- you know, we could use the Economist McDonald Index or some variant of that, but...

DR. CHESNES: Any other questions?

(No response.)

DR. CHESNES: All right. Thank you very much.

(Applause.)

DR. CHESNES: So, now we are going to have our first panel session on the economics of consumer financial protection, chaired by Janis Pappalardo.

PANEL SESSION ONE:

ECONOMICS OF CONSUMER FINANCIAL PROTECTION

DR. PAPPALARDO: Welcome this morning. We're really pleased to have the panel to address the economics of consumer financial protection. The overarching theme of today's panel is going to be to try to give those of you who are not familiar with consumer protection questions a sense of the many interesting questions -- in other words, welcome to the wild world of consumer protection -- in the consumer financial area.

We've tried to have a general theme for the panel today, which is access and price of credit, because consumer protection policies broadly affect access to credit and the price of credit. What we'd like to do is have our speakers talk in some detail about areas in which they are expert to give you a sense of the many rich questions that still remain to be answered in the area. It's a fabulous panel.

We have Dan Becker from the Federal Trade Commission, Jesse Leary from the Consumer Financial Protection Bureau, and Jeremy Tobacman from the University of Pennsylvania, The Wharton School.

I am going to present in sort of a different order, one that would make a little bit more sense given

the topics that they are going to talk about. Jeremy is going to talk first about access to alternative financial products, particularly, I think, payday loans. Second, Jesse will talk about the role of credit scores and credit reports and pricing and availability of credit. And, finally, Dan Becker will talk about the role of debt collection in the credit market.

I will make a few comments at the end, and then we'll open it up for questions. And I'll ask you to hold your particular questions, unless they're really urgent questions of clarification, until the end of all three speakers. Thank you.

Jeremy?

DR. TOBACMAN: Okay. Thanks, Jan, for inviting me to be here, and thank you all for being here as well.

So, I was asked to be quite broad in talking about the economics of consumer financial protection in the realm of payday lending, and I'm also charged with being quite brief. So, this leaves me at some risk of saying almost nothing at all. So, to compensate as best I can for that, I'm going to err in what I say on the side of being blunt and provocative. So -- but I hope you won't take anything I say as the final word because of this caveat.

So, because this is an industry that's grown a

great deal over the last 20 years and now may be shrinking, but regardless, I want to try to make sure that we're all on the same page with a couple basic facts about payday lending. So, first, the typical payday loan scenario involves a borrower going into a storefront -- think, you know, in an urban center or a suburban strip mall type area -- receiving $300 in cash following a very rapid-fire underwriting process, and then handing over a paper check dated two weeks later for $354, 18 percent more than the amount of the loan that they receive.

So, there's variation, of course. Typical loan sizes are between zero and $500 or $100 and $500. There are some states where the exemption from small loan -- the small loan exemption from standard usury laws has set a -- has expanded the payday lending business to loans of up to a thousand dollars, but this is sort of a typical scenario to have in mind.

There are different estimates on the size of the market, but there was a payday lending question added to the SCF in the most recent wave. The estimates range from between 2 to 10 percent of U.S. households taking out a payday loan each year. And these are households that do have an account on which a repayment can be drawn. So, they're not the -- they're not unbanked

households and they're not the poorest of the poor.

The volume numbers that I have are a bit outdated, but the numbers from an investment bank that advises the industry extensively are that payday lending, that the annual flow of these loans is about $40 billion. So, just to give a sense of the market a little bit more broadly, physical payday lending locations offer lots of other services, which are sometimes lumped together as "fringe" financial services, including check-cashing, money orders, and pawn loans.

The online market share has been growing extremely fast and may even now constitute a majority of the business in the U.S. I'm not sure about this. If anybody has recent market stats, I'd be curious. And I think that, you know, by all accounts, this seems like a pretty competitive market. Concentration numbers are -- seem to be -- seem to be fairly low. Entry costs are pretty trivial, because these are just sort of ordinary storefronts, and a third-party credit bureau, called Teletrack, serves this industry and other subprime lenders, which -- and reduces the informational advantages of any incumbents.

So, I think, you know, this is sort of the broad picture of what the market looks like. I haven't told

you anything that speaks to consumer financial protection issues yet. So, I'm going to try to get this now, but this is sort of a starting point for the picture.

Before I talk about possible rationales for some intervention in the market, however, I want to mention the evidence that we've seen on just straight-up impacts of access to and use of payday loans. There are a bunch of papers in this industry on this question looking at a range of outcome variables, from bankruptcy to bounced checks, and I think that the evidence across these is mixed. I'm a little bit biased, because my name is in the list. I'm not going to go through this in great detail, because John Caskey has a graciously devastating review of all of the work on these questions. So, that's where you should turn.

Instead, what I want to focus on are biases in decision-making and the ways that they may have consequences for consumer behavior and the structure of the market for payday loans, and there are three particular biases that I want to focus on.

The first is extreme impatience, especially in the short run. The second is overoptimism, and by this, I mean potential overoptimism about a variety of variables. And the third is low levels of financial

literacy or financial awareness, and I am going to try to focus the most on that area, in part because of the expertise in this building in disclosures, where I think that there's some interesting work that speaks to the way that disclosure policy might evolve.

Okay. So, the first bias, extreme impatience. So, there's an abundance of evidence across hundreds of domains, thousands of papers, that consumers exhibit extreme -- that consumers often exhibit high annualized discount rates and also that discount rates in the real world, in the field, seem to be higher in the short term than in the long term. And there are a variety of models that have been used to capture these notions -- models of temptation, hyperbolic discounting, self-control -- that it seems natural to explore as a possible explanation or a possible relevant fact in payday borrowing, because the interest rates are rather high, higher than most other forms of credit in the U.S.

So, one of the difficulties in taking these models to data in the payday lending context is that it's often hard to separately identify shocks to income or consumption needs from discounting unless there are other choices that consumers are making simultaneously. The signature implication of models of self-control is a demand for commitment, and when we're just observing

behavior about borrowing or not, then we can't tell if people are purely impatient or if they have self-control problems.

One implication that I think is important is that we would expect that sophisticated hyperbolics would default quickly, because by doing so, they would commit themselves to not borrow subsequently on such high-interest credit. Now, and actually, this is effectively the test that we're able to run in one of our papers, and we find that actually the borrowers seem to delay default by a substantial amount of time, suggesting that sophisticated hyperbolic discounting is not what's going on in this market, primarily.

So, that brings us to overoptimism, and overoptimism can be parameterized in a variety of ways, overoptimism about income or future shocks or default costs or overoptimism, also, about your future degree of self-control.

So, the two papers that should be on everybody's radar screen in trying to think about how overoptimism operates and matters for a variety of consumer markets beyond just consumer credit are the Dellavigna and Malmendier paper on contract design and self-control, which demonstrates equilibrium contracts, two-part tariffs, when -- in cases of monopoly, competition,

sophisticated hyperbolic discounting, and naive hyperbolic discounting.

And the key results in this paper -- key result -- combination of results are that in the case with perfect competition, there are no welfare losses from sophisticated hyperbolic discounting because the initial inducement to enter the contract for a temptation good compensates for the follow-up payments that the consumer will pay. Think of this as the companion ticket you get when you sign up for a credit card that has a high post teaser interest rate.

But there are welfare consequences under perfect competition when consumers are naive, because essentially too many consumers sign up. The initial payments are sufficient inducement to get them to sign up, but because they underestimate the degree to which they're going to borrow on their credit cards after the teaser rates expire.

So, Gabaix and Laibson, at the start of their paper, replicates with different language the results in Dellavigna and Malmendier, but then extends and shows how shrouding of delayed characteristics of products may persist in equilibrium, and by introducing a clever structure for advertising in that context.

So, the key or one key conclusion from these

papers is that in the presence of naivete, competition does not or at least may not restore efficiency, depending on how you structure the information evolution process, and, in fact, the market may -- may persist in sustaining characteristics that look like naivete or sustained overoptimism.

And we have a variety of contexts where there's overoptimism in consumer financial decision-making. The last of these papers, by Ausubel, is the best example.

Okay. So, our results on payday borrowing and overoptimism come from the following logical chain: The typical borrower at a large lender borrows repeatedly and eventually defaults at some point within one year of their first loan. Conditional on default, they have already paid 90 percent of the original loan's principal in interest, and so the question then is, if they rationally forecast this probability of eventual default, why wouldn't they have defaulted immediately?

And so interest payments preserve the option to borrow subsequently, but the estimated structural model in this paper implies that the value of this option is small. If they're already discounting enough to borrow in the first place at 18 percent for two weeks, then the option value, which you would think would emerge over an extended period of time, just can't add up to much.

So, this is what I said about procrastination on defaults or delayed defaults as being inconsistent with sophisticated hyperbolic discounting but naturally explained by naive hyperbolic discounting, okay?

So, my time is coming to a close. Let me try to say a couple quick things about low financial literacy, which also affects how we think about this market. So, I think that the most powerful results or evidence on this is the direct work by Lusardi, Mitchell, and Tufano in separate papers, just documenting very, very low levels of financial sophistication and sophistication about debt in the U.S. population.

The indirect evidence I think comes in -- we have in two potential forms. One is a finding from some of my own work that people who borrow on payday loans oftentimes have substantial access to other forms of liquidity, amounting to annual losses from the failure to arbitrage of about $150 per year. There are selected samples, but this -- and this is not evidence of violations of a vervillick (phonetic) preference, but it suggests that there are gains on the table.

And then the other context is this randomized field experiment on information disclosure by Marianne Bertrand and Adair Morris, and they had -- and I think that actually this is a wonderful paper with results

that we should all try to absorb and internalize.

The reason it's such a good paper is because they brought to bear a huge variety of very powerful insights from the psychology and economics literature to try to make disclosures that would cause consumers to make better decisions in the realm of payday lending. They gave consumers aggregate information. They helped consumers make calculations and comparisons, and they made all of this information extremely salient by putting the new information on the physical envelope in which people got their $300 in cash after getting their payday loans.

That's something that they were going to take home. It was in their hands. They could look at it later. It was an unavoidable reality of this information that was available to them. So, overall, this is a wisely designed, carefully implemented, informative experiment about the effect of disclosures.

Now, so, what are the findings? Well, two of the three treatments had small and insignificant effects. There were some significant effects in subsamples but not a whole lot of information there, I don't think. The punch line that comes out in terms of statistical significance is the effect of their dollar treatment, and I should go back.

So, this was disclosed, how fees would add up, would accumulate for up to three months, and so this was just reported on the back of the envelope. What they find is a significant effect that this treatment reducing the fraction of subsequent pay periods on which people borrowed from 54.2 to 48.7 percent, statistically significant. This was an effect over the next four months.

So, how do we think about the size of this effect? Well, in some sense, it's huge, right? This can be implemented for zero marginal cost. It's a form of benign, limited, libertarian paternalism. It's just, you know, like other disclosures. It's a very -- in some ways, a very gracious intervention.

In some ways, this seems like a minuscule effect. After being confronted very, very baldly with the fact that payday loan interest adds up quite dramatically, almost half of the population continues to borrow in every pay cycle following this intervention. And so I think that the -- and the effect was actually very small on the people who were the heaviest borrowers.

So, I think, you know, this suggests to me that the role of disclosures can be expanded, but to the extent that low financial literacy is influencing or is

influencing consumer decision-making, then perhaps we need much broader-based interventions to overcome it.

Okay, so opportunities and questions, very, very quickly. I think that, you know, when we talk about overoptimism, when we talk about low financial literacy, we're very interested in how quickly and how much and how do consumers learn to make better decisions, what interventions can be constructed to do that.

In terms of data, account-level data are potentially available and can be used to try to understand the role of biases and the effects of access. In addition, financial institution partners and especially credit unions have been open about sharing data and trying to learn more about how this market works and how decision-making in this market works, what are the effects of access and use of these high interest rate products that are so common.

There have been lots of regulatory changes. States have been legislating. CFPB is going to have substantial authority, once they have a director, over payday lending. And then there is all this expertise in this building about Truth in Lending oversight that I think can also inform subsequent decisions here.

Okay. So, I may have been too bold, but at least here are some things on the table. Thank you.

(Applause.)

DR. LEARY: All right. So, I'll handle the transition. That looks like it should be it, yeah. Let's see. Okay, I don't want everybody to see my notes.

Okay. So, I'm Jesse Leary. I am an economist at the Consumer Financial Protection Bureau. Before going to the Bureau, I was here for 10, 11 years, working on a wide range of consumer protection matters, including consumer credit reporting and credit scoring.

A quick disclaimer: The views in the talk are my own, not necessarily those of CFPB, the Treasury, their staffs. The other disclaimer about the terms I'm going to use, the way everyone talks about credit reporting and credit scoring is a different set of terms than what's in the law, and so just -- I don't want to have to, like, be going back and forth all the way through. I see at least one FTC attorney in the room, so I wanted to, you know, point out that I do know what the right terms are, but we don't have time to do it.

So, what I'm going to talk about -- what I'm going to do is give a very brief overview of what credit reporting is, what credit scoring is. I don't know how much everybody knows, but I'll try to just give a very basic foundation. And then I'll talk about potential

consumer protection problems in each of these areas and about reports and studies being done here at the Commission and at the Bureau in each of these areas.

When I was here, I worked on some of the studies. Others, I didn't really work as much on directly, but I'll talk about those as well. So, I'll talk about credit reporting first and then credit scoring.

So, briefly, what is credit reporting? Credit reporting is collecting information from creditors, from public records, and lots of other sources, and compiling that into credit reports. There are the big three firms that do this, TransUnion, Equifax, and Experian. There are lots and lots of specialized firms pulling various other kinds of data. Jeremy was referencing Teletrack that does this for the payday -- sort of in the payday space, but typically, we're talking about the big three that put together credit reports that try to be fairly comprehensive. So, collecting this data, compiling these reports about individuals.

The reports are used very widely. They are used, obviously, by lenders to make lending decisions, but they are also used by insurance companies. They are used by utilities. When you're setting up a gas line, they might pull your credit report to see whether to

require you to put down a deposit. They're used by landlords. So, I actually used to -- I actually used to be -- I used to have a -- be a landlord, and so I would pull credit reports on potential renters, as this is a fairly common use.

And they are also used by employers. This is a fairly controversial area, where there's -- I guess there's some -- there's some dispute about whether it should -- about whether an employer needs a very specific reason to be thinking about -- to be worried about financial issues with an employee or whether employers can pull it as a general sort of character check on people.

One other thing about reports, in addition to gathering data from these places and putting that on the report, the credit reporting agencies keep track when someone requests your report, because that can also be an important indicator of credit risk.

So, what are credit scores? So, a credit score is a numerical summary of the -- a credit score takes the data in your credit report and turns it into a number that summarizes that information, and typically, it's done to indicate the relative risk that you'll default on a loan. There are -- the folks who know anything about it tend to -- what they tend to know is

the FICO score. So, FICO is the dominant seller of third-party credit scores. They invented credit scoring 60 years ago and have been the dominant player ever since.

Recently, the credit bureaus have formed a joint venture called Vantage Score, which does the same sort of thing. It generates scores based on credit reports. There are also an enormous number of proprietary scores that are in use by individual lenders. So, any -- from talking to folks in the industry, what you hear is essentially any sizeable lender is not just pulling the FICO score or Vantage Score. They are actually building their own models internally, pulling credit report information, sort of raw credit report information, and generating their own scores.

And then there are also credit risk insurance scores, which are very similar to credit scores in that they -- you know, they take data from a credit report and they run it through an algorithm which will generate a score, but what they are predicting is something about your likely future claims on a homeowner's policy or an auto policy.

And there are also, just to make it more complicated, there are scores used by individual lenders. They are -- some proprietary scores simply

take data from credit reports. Others will take that data, combine it with data from an application, for example, and come up with a summary of your likely riskiness that's based on a broader set of information than just what's in your credit report.

So, moving on to potential consumer protection issues and starting off with credit reporting. So, the customers for credit reports are typically lenders or, you know, the other customers we talked about, but lenders are the biggest users of credit reports. So, consumers benefit enormously from the credit reporting system. I mean, it enables you to go to a bank and sort of credibly convey your credit quality.

But the consumers are not the ones who are actually paying the credit bureaus, and so when credit bureaus are deciding how are they going to build -- how are they going to structure their systems, how are they going to -- how are they going to compile information, and how are they going to resolve uncertainty about is this piece of information accurate, does this piece of information relate to this consumer, absent any policy intervention, what we'd expect them to be most concerned about is how does -- what are the demands of -- what the incentives are of our customers and how do we translate that into decisions we are going to make about how we

construct credit reports.

So, lenders obviously care about accuracy. They want -- you know, they want to make -- they want to be making decisions based on accurate information, but there is going to be some inaccuracy in the process, and the costs to lenders of that inaccuracy are not symmetric. So, if I think someone is a good risk and it turns out they're a bad risk and I make them a loan, that is much worse for me as a lender than if I think someone's a bad risk when really they're a good risk and I forgo a loan that would have been a good loan. So, there's a concern that lenders will -- that the credit reporting agencies will tend to be overly inclusive of negative information when there's some uncertainty about the accuracy of that information.

In addition, there are costs to consumers of inaccuracy. If you're not able to get a loan, you're not able to rent an apartment, not able to get a job, because of inaccurate information on your credit report, that -- you know, there are costs to the lender of the -- sort of the mistake, but there also can be substantial costs to the consumer, and that's likely not internal to the decision about the credit -- that the credit bureau is making when they're compiling reports.

So, absent regulation, we would expect to see

too much inaccuracy and inaccuracy that tends to go in a certain direction, and as a result, there's the Fair Credit Reporting Act that governs the behavior of credit bureaus and of folks who provide information and users of credit report information.

So, the FTC is working on a study of credit report accuracy. So, the big question that comes out of that is just how -- the headline question is just how inaccurate are credit reports. So, the FTC is working on a major study of credit report accuracy. There have been two pilot studies over the last six years or so to develop a methodology that should be effective at identifying at least some of the most important sources of inaccuracy. I didn't work directly on this study.

I'm looking to see if -- I don't see Peter Vander Nat in the room. He's been the driving force on this study. So, it's too bad he's not here, but on the other hand, if I screw up, it's less likely to get caught.

So, what they're doing now -- and I believe this is in the field currently -- is they're reviewing credit reports with a large number of consumers, so having someone who's experienced in reviewing credit reports go through a report with a consumer to help the consumer understand what's there and help the consumer identify

information that the consumer believes to be inaccurate.

They will be identifying -- so, there's lots and lots of information on a credit report. Some of it might be inaccurate and not matter, but there -- but -- whereas other things could have a real impact on the apparent creditworthiness. So, the focus will be trying to identify things that would actually have a material impact on the creditworthiness as portrayed by the credit file.

And then consumers will be encouraged and assisted to dispute those items that they believe to be inaccurate with the credit bureaus, and that's going to be an important way of trying to learn as much as possible about whether these things are inaccurate. If you just take -- there's obviously some risk that consumers -- if you just ask them what's accurate, what's inaccurate, they are likely to tell you negative information is inaccurate. So, this is a way of trying to -- this won't -- this doesn't get you to truth, with a capital T, but it gets you more information about whether the information is likely inaccurate. And that study is going to come out next -- well, it's due next December. I don't know whether they're ahead of schedule or not.

So, the two primary concerns about credit

scores, one is transparency. Credit scores can play an important role in consumers' financial lives, but consumers don't know a heck of a lot about it. That leads to anxiety, concern on the part of consumers, on the part of policy-makers. I think there's a lot of scope for figuring out what consumers really do need to know and what they don't need to know, so what would be helpful information, what would not be helpful information. But there's a lot of -- you know, there's a lot of concern in the area.

In addition, another important fact about credit scores is that the scores of racial and ethnic minorities, particularly African-Americans and Hispanics, tend to be much lower than those of non-Hispanic whites, as well as the scores of recent immigrants, and to some extent, by income, although income is not -- the difference by income is not that dramatic.

So, I am going to take these in reverse order. So, this shows -- this is from the FTC study of credit-based insurance scores, and it shows the distribution of scores by race and ethnicity. The yellow line -- I should have a laser here. So, the yellow line up here, this shows that -- this is showing score decile, and this shows the share of consumers by

race or ethnicity that are in each score decile.

So, more than a quarter of African-Americans are in the bottom decile of credit-based insurance scores; nearly 20 percent of Hispanics are in the bottom decile of credit-based insurance scores; and then much lower down here, in the high score range. As I said, this is credit-based insurance scores from the FTC study.

The Fed did an analogous study of credit scores used in credit markets that was released about the same time and had nearly identical results. So, you know, this sort of raises some obvious equity concerns. The FTC study of insurance scores showed that using scores, it would be likely to have a pretty substantial impact on the premiums paid by African-American and Hispanic drivers -- in this case, it was auto insurance -- but that very little of the relationship between score and risk is omitted variable bias.

That is, scores are not powerful because they're correlated with race or little of their power comes from a correlation with race, and so the -- there's a -- sort of the first order of concern of if you allow the use of these scores, there will be distributional effects. There was a parallel concern that the only reason scores matter is because they're sort of acting as a proxy for race. I think the FTC study showed that that's a very

little -- that that's a -- in my view, a -- you know, a pretty small component of what's going on with scores in the insurance market.

Similarly, in the credit markets, the Fed study showed similar results, you know, that scores are predictive -- scores -- there are big differences in scores across groups, but scores are quite predictive of risk for all groups. And the one finding of their study that I think raises some issues is that for relatively recent immigrants, people who immigrate to this country as adults, their credit reports tend to look like younger people.

Younger people are risky, and so immigrants are safer than their scores would -- on average, are safer than their scores would tend to indicate. So, I think that that's one area where there's a real potential for concern. It's hard to know what to possibly do about that.

We're actually working on another study involving the potential use of remittance data, which is data on people sending money to foreign countries, to see whether that could be incorporated into credit scoring and what role that could play in possibly addressing this issue of sort of the underprediction of creditworthiness of immigrants, but that -- and

that's -- but there's some real sort of -- it's not clear that's a fruitful avenue. We've been sort of asked to explore whether that's a fruitful avenue.

So, quickly, we're doing -- and I should say, all the studies I'm talking about here are all congressionally mandated studies. These are all things that Congress has asked the FTC or the CFPB to do, which I think speaks to sort of the level of policy concern about these topics.

We're doing a study now or we released a study in -- we released an initial study in the summer of the differences between scores that consumers are able to buy -- if a consumer wants to know their credit score, they can go into the market and buy a credit score. Those don't tend to be scores that are actually being sold to creditors. So, that raises a concern that are people buying something that's useful to them and are they being somehow led astray with the use of these scores.

So, if consumers -- if you do buy a score and it gives you a very different -- it gives you a very different message about your creditworthiness, potentially that could lead you to sort of apply for the wrong kind of credit. If you think you are riskier than the market thinks you are, you might apply for or accept

credit on worse terms than you could actually obtain.

So, we released a study in July that sort of lays out the background information and raises this as a potential issue, but right now, we're working on a fairly simple analysis of a very rich data set where we're getting -- we have credit reports from each of the three bureaus on 200,000 individuals -- well, there are different people across the bureaus, but we have 200,000 people from each bureau, and we have the FICO scores, which are the scores that are most widely used by creditors, and we also have the scores that are sold to consumers, and we will be able to see how different they are.

You know, if people are buying -- going out and buying a score on the Web, are they going to -- is it going to tell them -- is it going to be sort of informative about their -- about their creditworthiness and how they will be concerned -- how they will be perceived by lenders? And the big unknown, I think, in this area is, how do consumers actually use this information when they get it? Does it actually impact their decision-making? And what will be most useful for them to be learning about these issues?

All right, and I've run over a bit, but that's it. Thanks.

(Applause.)

DR. BECKER: I'm going to discuss some of the challenges of policy in the debt collection industry. This is an area that the FTC is in the middle of a study right now, but what I talk about today is not so much going to be about our ongoing study, but it's going to be more about what are the challenges and what are the -- I'm going to describe some of what I believe are the open and interesting research questions.

So, the debt collection industry, from an economic point of view, is a very interesting industry, because it has sort of all the problems that we sometimes think might arise in any given market. So, we have multiple sources of moral hazard; we have multiple sources of adverse selection; we have consumers who are frequently facing a budget constraint that's binding in a particularly real way; and it has this unique feature that the firms are spending money when they call consumers to collect explicitly to more or less bother the consumer that they're interacting with, because that's basically how you collect. And so it has lots of unique features. Some of them, in some sense, can be sort of problematic, and it also is a really large industry.

The face value of debts that are sent to

collections each year is about $110 billion. So, that includes credit card debts, auto debts, medical debts, debts to state and local governments, and then those can go -- the model under which those are sent to collections, there can be either that they are simply sold to a collector who can then collect on them as the owner of the debt, or some of them go into contingency collection, and then the collector will collect as an agent of the debtholder.

But both -- and when these things are sold, they're typically sold, depending on the type of debt, for about 5 cents on the dollar. So, you can think about market size, if you multiply this by 5 cents on a dollar, we're still talking about $5 billion, which is real money. And because of sort of these intricacies of the market and the fact that the market is really big, this has over the years been the largest source, in terms of industries, of complaints that the FTC receives.

So, the market's important not only because of its size, but I am going to argue that it has effects that range -- that are pretty far-ranging throughout the economy. So, the reason for that is that having an effective collection system lowers the cost of credit to consumers. That is -- the mechanism for that is that if

a creditor is able -- knows they're able to sell their debts, they're able to recoup some money that they otherwise wouldn't have received, and that increases their willingness to offer credit or increases the supply of credit.

And then what I think is probably a bigger effect is that people know that there's an effective collection system, and so there's this deterrent effect. Roughly 90 percent of accounts never go into collections, and the reason for that is that people know that there's an effective collection system out there.

So, we have increased access to credit, and that facilitates all sorts of household investment, whether that be education, the ability to pay sudden medical bills, the ability to invest in a car so that someone can get to work, and there's actually literature that says that access to credit and trust in contracting has a large impact on economic growth more broadly. So, I think it's pretty clear that this is a very important market.

So, I mentioned that the effectiveness of the collection system has some effect on the price of credit. I just want to talk a little bit more about the mechanism for that. So, when we have -- when money is collected or when creditors know that money will be

collected, that raises the cost of these debts when they're resold, and that, as I said, increases the supply of credit. And that relationship of -- so, that causes creditors to be more willing to make loans, and we see that a lower interest rate is how that gets -- money gets transferred from people who otherwise wouldn't pay their debts to those who do pay their debts but now pay a lower interest rate.

If the debt collection market and the credit market are both perfectly competitive, that can be close to a one-to-one pass-through or transfer from those who -- those nonpayers to the payers, and then if the markets are less competitive, then that pass-through can be on less than a dollar-for-dollar basis. And this is what the -- the structure of the market is one of the questions that we are trying to address in this ongoing FTC study.

The other reason that this pass-through might not be on dollar-for-dollar basis is that you actually have to pay people to sit there on the phone and call the debtors, and that has some cost. So, those are the two sources of the sort of loss that keep that from being a pass-through, but otherwise, it's more or less a transfer.

So, frequently, in that kind of transfer, there

isn't really an efficiency aspect to that, but I am going to argue that when we think about collection policies, there is a big efficiency aspect to this or this is potentially a dead weight loss.

So, the example I am going to give is if you imagine a policy that restricts how collectors can collect, and as a result, they collect less; as a result, the equilibrium interest rate increases from, let's say, 10 percent to 15 percent. You have lots of inframarginal consumers who would have had some benefit if they could take out a loan at 11 percent or 12 percent or 13 percent, and they're made worse off, and the amount they're made worse off isn't being transferred to anyone else. There's no one -- there's no direct benefit from the fact that these people have been effectively priced out of the market. So, that is some efficiency loss or dead weight loss from restricting collections too much.

At the same time, we like to restrict collections some, because collections is inherently sort of a painful process to those who are involved. I think of it as a series of sticks but no carrot. So, a collector calls you, they say, "We're going to bother you and we're going to keep bothering you until you pay us," and that's just a direct harm that doesn't really

have an intrinsic benefit to anyone else.

And to mitigate that, there's a law called the Fair Debt Collection Practices Act. The Fair Debt Collection Practices Act restricts certain collection techniques. You can't call someone and threaten violence; you can't call them in the middle of the night and wake them up; you can't call their friends, their neighbors, and their families and embarrass them that they're not paying some debt. And, you know, there's an efficiency basis for that. We don't want collectors imposing some large direct harm on consumers.

So, you can imagine sort of a frontier of what collection techniques we allow versus what prices are associated with that, and I like to think of the loan as being a bundle of that price and the collection protections.

Jesse actually introduced me to an analogy that I kind of like, that he was making a loan to me, and he couldn't pull the money out of my wallet, or perhaps he can't see how much money I had. We might sign a contract where if I don't repay him, he gets to punch me in the face. And that's going to make me more likely to repay him, and as a result, the interest rate that we agree to will be lower.

And you can even imagine if we were to form a

contract and we could potentially contract how hard he hits me in the face, even though I -- at least I hope he doesn't get any intrinsic benefit from punching me in the face, he might want a contract that he can punch me in the face harder, and I might agree to it because I can get a lower price, because there's some assurance there that I'm actually going to pay him back.

So, there is some optimal bundle in both this example and you can imagine an optimal bundle that people would choose from somewhere on that frontier in terms of what collection techniques would be allowed. Empirically, we don't see any variation in this. We don't see people contracting for what techniques would be allowed, and I think there are two possible answers.

The one that I think is probably less likely is it could be that the things we would contract to are already prohibited by the FDCPA; that is, if we had a choice, we would allow these collection techniques that are more aggressive, and that's already prohibited by law. I think that's not it.

I think it's more likely to be adverse selection. If I go to a bank and say, "I want to sign a contract with you, I'm even willing to agree to a higher interest rate if you promise not to be very aggressive if I fail to pay you back," that kind of tells you

something about whether I intend to pay you back. And then, similarly, it's sort of just a messy process to contract for, and I would guess that people probably don't even -- that most consumers don't know what the current credit collections are and they don't know what the laws are.

Okay. So, I said I will get to the research frontier. I think the questions, some of which we're trying to address but quite a few are going to remain unanswered even after study has been published, the first is, what is that relationship between the amount that's collected and the equilibrium interest rate? This is something that I said we're trying to get some insight into through some insight from theory about how it's related to industrial structure, and then we're actually looking into the industrial structure.

The other part of that, from the collection side, is how do the laws, in terms of what's allowed -- what we allow people to collect, how do those affect the amounts collected?

On the consumer side, the basic question is, are consumers even aware of what protections they have under the law? And then the second question is, if they are or even if they aren't and we were to educate them, is there some moral hazard issue there where they're going

to use these protections in a way that they weren't intended to default on loans that we think they should not be defaulting on? So, that's the collectors and the borrowers; that's really the two parts of the market.

I think there's some big-picture questions that we could address potentially even without being able to pin down parts of those two separate sides of the market. The first is, given how many moving parts there are, is there some way that we could find the optimal policy, even without figuring out what the moving parts are?

So, I'm sort of thinking of an efficient statistics strategy, if you could imagine moving along that frontier. Is there something that we might observe that would tell us what the welfare-maximizing point on the price versus protections frontier?

Something that sounds appealing, though isn't exactly it, is quantity. So, the set of prices and protections that maximize quantity don't necessarily maximize welfare, though I haven't thought of anything else. If someone could do something clever there, there would be a lot of value in that.

What I was sort of thinking of as the million-dollar question, though it's worth much more to us, is how many mistakes are there? The complaints we

get very frequently are that people claim that the collection companies are trying to collect the wrong amount or they're collecting from the wrong person. There's been a change in technologies in how these debts are transferred, but consumer advocates say that there are a lot of mistakes, and even collectors will say that mistakes happen.

So, the real question is, how many mistakes are there in the collections process? And then, how would the -- how would different technologies affect that? And there may even be some scope for policy in encouraging the adoption of new technologies if there's some sort of coordination problem in adopting those technologies.

And then, lastly, just thinking of -- if you were to really zoom out and think about the game that collectors and borrowers are playing here, and debtors are playing here, it's this game of attrition. There's sort of a painful game of attrition. You have got someone calling you, potentially on a daily basis, and bothering you and perhaps making different threats, hopefully legal threats or threats that are legal to make, and at the same time, the collection company is spending money bothering these people.

So, it's really a mutually painful game of

attrition. To the extent that we can either make this process smoother and make it less painful for people, for both sides of the market, that would be beneficial. To the extent that we can at least root it into something that has other benefits, whether that be making sure that its effect in terms of credit reporting is reasonable, so that the credit market in the first place works better, I think that zooming out from all these little questions and thinking about how can we improve this game or how can we improve this market, given that it's $110 billion market that has implications elsewhere, I think the benefit to that -- the benefit of improving this, to addressing any of these questions, would really be huge.

(Applause.)

DR. PAPPALARDO: Okay. We're a little bit late for time. I would like to take a few moments to just make some observations about some of the big-picture questions that were raised today.

So, the three topics that were discussed all have to do with availability and price of credit, which is sort of fundamental to the ability of a household to smooth their consumption over their lifetime and to invest in education and other things that could improve their life as time goes on. So, these are big

questions.

I hope that the idea that you got from today's panel is that consumer protection policy offers everything an economist could ever want. We have all kinds of property right issues, and these are fundamental questions that extend beyond just financial protection. So, who owns information about me?

Since there are transaction costs to information, this market for information about me really raises lots of interesting questions for economists in the area of privacy and privacy protection, as well as in, more narrowly, the financial protection area.

As we've heard today, we have moral hazard issues. We have lots of imperfect information issues, lots of problems that happen because it costs time and money to search for information and then to comprehend and understand, to transform information on the page to something that people can actually use.

We have equity efficiency trade-offs, like in the debt collection area. We have all kinds of externalities. We have all kinds of competition questions, information asymmetries, risk and uncertainty. That is so fundamental to trying to understand this market and so important to understand in thinking about consumer protection.

I know we are going to have a mortgage panel following this panel, but one thing that you have to keep in mind when you look at the mortgage market is when people buy a house, they're often buying a consumption bundle as well as an investment bundle, and to the extent there are a lot of people who would like to reduce risk, you have to ask yourself, what does that mean in the long run for people's ability to move ahead?

We have liability questions. We have all kinds of parties involved in transactions. Who should be held liable? What's efficient? What's the efficient strategy?

We also have all kinds of -- consumer protection also offers everything a behavioral economist would want. We have all kinds of optimism issues and patience issues, loss aversion issues.

Different theories of consumer behavior can lead to different consumer policy recommendations, and I think it's really fundamental for people to understand, who are not in the consumer protection area, that I think in this day and age, behavioral theories tend to dominate the discussion in policy circles, okay?

You may not always see it in print, in academic papers, or you might sort of get a sense of that as time goes on, but I think that that's the fundamental

prevailing paradigm these days, especially amongst a lot of attorneys and noneconomists who really have the policy power to make changes.

So, suppose you have people and you see that they're not maximizing income. Some people interpret that as saying that people are behaving irrationally; they're leaving money on the table. Another interpretation is that people are rationally maximizing utility -- not just income -- subject to wealth constraints, time constraints, and household production constraints.

I think it's really important in this area that people understand objectives and constraints, to distinguish between these two types of interpretations, because different types of interpretations have fundamentally different policy recommendations that are attached to them.

An example that came up earlier today was FTC mortgage research. We found in our research that consumers might think that they have truthful, accurate information on government-mandated forms, but that information is put on the page in a potentially misleading way. So, people who seem to think that they're rationally responding to information given to them by the government may make bad policy choices for

themselves. So, we have to distinguish between bad decisions because the information environment is flawed and bad decisions because consumers are somehow behaving irrationally.

It's a huge debate right now. There are many people, based on behavioral theories, who like to eliminate or limit access to high-cost credit products and nudge people to what is considered, in the policymaker's view, the preferred choice. The jury, I think, is still out on behavioral versus microeconomics models of consumer behavior and which would lead to the greatest welfare for society.

As an example, I would recommend looking at a paper by Greg Elliehausen, where he talks about what behavioral views and more traditional views of consumer behavior are showing. He says, "At this time, neither existing behavioral evidence nor conventional economic evidence supports a general conclusion that consumers' credit decisions are not rational or that markets do not work reasonably well."

In the payday loan area, as was mentioned earlier, there are different papers with different findings. One paper finds that the implication is that access to finance can be welfare-improving even at 400 percent APR. We have another paper where the results

suggest that restricting access harmed Oregon respondents, at least in the short term, by hindering productive consumption, smoothing, and/or investment.

So, my message to you is join the debate. Clarify the theories of consumer behavior and test alternative theories. What we need are more tests of the alternative theories so we can understand which ones are more predictive to help policy-makers. We need more empirical research.

I think this paper by John Caskey was mentioned earlier today in the payday loan area. His review suggests that there's an important public policy question for empirically oriented economists to tackle, that we don't know the answer yet. And we also need to clarify the microeconomic models of consumer behavior.

So, I've been at the FTC for over 25 years. I have, like, moss growing on my back. And we have a sort of a way of trying to combine models from the economics of information literature and other parts of microeconomics to approach consumer policy, but I think we haven't been really good at sort of clarifying how these models fit together, and we're trying to make more progress on that front.

So, we have two economists, Dan Becker and Doug Smith, who are working on our disclosure project, and we

have asked them to try to tie some of these models together in one coherent sort of set of findings.

And their preliminary findings, from a model of rationally inattentive consumers, emphasizes the role of information and consumer choice consistent with the FTC approach, but it ties it up more neatly. And what they're finding so far is that consumers are more likely to use information and their welfare will improve if information is less costly, and by less costly, also easier to process and to understand.

I think this is where there's some overlap between the psychologists, the marketing researchers, and the behavioral economists, that we're really trying to understand what people get when they read or search for information. Consumers are more likely to use information when they think their use is important. Information use is endogenous. Disclosures are likely to be more efficient than mandates -- direct product regulation -- if consumers have heterogenous preferences and individual consumers have accurate beliefs about whether the disclosure will be worth their time.

In a world where we have limited access to credit in recent years, the question I ask for you, in light of this situation, would people be better off with regulations that reduced credit options or not? This is

the level of debate now in the consumer financial protection area. It's a debate that's ongoing, and I ask you to join the debate.

And if you have any questions, we have just about maybe two minutes.

(No response.)

DR. PAPPALARDO: Well, the debate will continue during the break, then. Thank you.

(Applause.)

DR. CHESNES: So, we will take a 15-minute coffee break and start back up at five after.

(Recess.)

DR. ADAMS: Okay, my name is Chris Adams. I'm an economist here at the FTC. Thank you all very much for coming to our conference, and I'm going to hand it over to Aviv to run our empirical analysis of mortgage markets session.

PAPER SESSION ONE:

EMPIRICAL ANALYSIS OF MORTGAGE MARKETS

DR. NEVO: So, in this session, we have three great papers that satisfy Joe's criteria. They are all written by really smart people, and they're all on very interesting topics, at least I think so. That doesn't mean that Joe has to agree with me.

I guess I should have the usual disclaimer, you know, what I say does not represent anyone's view, not commissioners, and blah-blah-blah.

So, the first paper is by Zahi Ben-David from Ohio State. I've been told that I'm in charge of keeping time, although I think Laura is looking over my shoulder. She should -- what Laura says, do.

So, I think we're going to have -- is it 20 minutes per paper, five minutes or so -- seven minutes for discussion, and then we'll have a few minutes for questions. We want to take them in between or for all papers together? Up to me? That's never a good thing. Let's try to take them -- you know, maybe a few in between, but we can also see if we have general questions later. Okay, so...

DR. BEN-DAVID: Okay. Thank you very much. Thank you very much for including this paper in the program.

So, the paper is about a consumer behavior in the real estate market, the residential real estate market, and I'm looking at the relation between leverage and willingness to pay. So, we've all heard a lot about the relation between prices and the availability of credit during the bubble years. There is both -- there are both empirical evidence and a theory that link the two.

Now, both in the empirical literature that exists and the theory, we often think about the market as an aggregate. It could be at the national level, it could be at the county level, it could be the zip code level, but we don't see a lot of microlevel evidence. So, it's not clear how more credit affects prices at the transaction level. Is it the seller who increases the price? Is it the buyer who is willing to pay more? How does this actually work?

So, what I have in this paper is data that is unique that includes both asking prices and transaction prices, in addition to mortgage data. This allows me to disentangle or have a closer view on what's going on, and I'll be able -- I'll try also to explain why we see this relation between leverage and willingness to pay. And I think that some of the results may be kind of -- have a different angle to what we think about --

generally about this relation.

So, the two main results, I would say, if I had to summarize them in two sentences is, first of all, there is very strong correlation between the willingness to pay and the leverage that buyers take or borrowers borrow, and it's especially around, as you will see in a few moments, around the full listing price. So, borrowers who are highly leveraged are much, much more likely to pay the full listing price than others. So, there is actually a discontinuity around the full listing price.

Now, when I look at a cross-section and try to explain this behavior, it seems that a lot of it is driven by buyer sophistication. There is a bit of financial constraints. There is a lot about real estate agents. Real estate fixed effects, actually, are very important in this relation, and I will have some little story that explains it. And there is some evidence, also, for an optimism of buyers with respect to a -- keeping a -- house prices, keeping a -- continuing to increase.

Okay. So, just before I show you the results, let's just think about several explanations why leverage and prices paid could be related. So, we could think about different stories. Here are a few of them. I

will not be able to test all of them, you know, kind of to the bone, but I will be able to say some things about different stories.

So, one story is about liquidity constraints. It's kind of a mechanical story. I want to buy a house that -- you know, my budget is kind of around $100,000, so I'm looking for houses between $95,000 and $105,000. I found a house that I wanted. It cost $105,000. So, it means that I need really to borrow the extra $5,000 between $100,000 and $105,000. And you see this relation -- you would see a mechanical relation between the amount that I'm willing to pay and my leverage.

Another story could be some behavioral or persuasion story. You know, I put little down payment, and it doesn't feel very heavy on my pocket, right? This is kind of the story that perhaps is going on in a paper by Levin and co-authors in the AR about the car market, right? The cars are sold for higher prices when there is no down payment or down payment is very small.

Another story could be optimism. If I don't buy this house now, next week, it's going to be more expensive. So, perhaps, you know, I should pay the full listing price, and, you know, I don't have this money, and I'm also a bit optimistic -- I'm confident that prices are going to keep going up, so I don't mind

borrowing that much.

And another story could be about moral hazard. So, moral hazard would be, you know, as a highly leveraged borrower, you know, I don't mind paying the full listing price, because I'm 100 percent financed. It's like an option. The market goes up, I win; the market goes down, the bank loses. In this case, in this particular story, we should see some constant effect over time, right, because it doesn't matter whether I'm in a period of rising prices or stable prices. I should behave in same -- I should have the same behavior.

All right. So, what kind of data do I use? So, I have the MLS data. So, these are all transactions that were conducted by real estate agents in the Chicago area from '94 until 2008. So, we're talking about approximately 770,000 transactions. And I have asking prices, I have time on the market, I know who the real estate agent was, and in addition, I link this database to a recorder of deeds. These are the entire universe of transactions in Cook County, in Illinois. It includes mortgage sizes, interest rates, and some foreclosure information.

In addition, I have income at the loan level reported in HMDA, and from Census tract, I take education information at the zip code level.

So, let me show you some stylized facts just to set the scene. So, my left-hand side variable, as in most regressions, is going to be whether you paid -- whether your price was higher than the listing price. It's going to be a dummy. Just for convenience, I multiplied by 100, so these are all less regressions. So, on the left-hand side, we have zero or 100. So, these numbers could be interpreted more or less like percentages.

So, what we see here, these are dummies for your leverage, LTV, loan divided by the price that you paid, and basically what you can see is that as your LTV increases, the likelihood of paying the full listing price increases by a lot. So, if you look at the entire period, we're talking about 13 percent more likely to pay the full listing price if you are at a leverage of 96 percent or above.

If you break it by periods, you see that it's strong in all periods, but it's especially strong during the bubble years. You know, I would note that there are very tight controls here, beyond the transaction controls, just because of space limits are not here. There are tax code interacted with quarter fixed effects.

So, this means that, you know, for every few

blocks, for every quarter, I have a fixed effect. There are, you know, about a dozen thousand fixed effects in these regressions. So, there is a very strong control for the geographic and time component here, and still you'll see that this effect is very strong, especially during the bubble years, but also -- or the peak of the bubble years, but also in other periods.

If we look at the -- if we try to isolate the effect of leverage after -- you know, after controlling and removing all the variations that I can control with, what we see in this chart is the relation between the percentage of transactions that paid the full listing price and leverage. So, you see that it's kind of more or less hovering around zero until 93 percent or so, and then it shoots up. So, it's really the guys with the leverage that is -- you know, goes from '97 onwards that tend to pay the full listing price.

If we look over time, we see here a three series. First of all, the guys with high leverage, 96 or above, we see that it's been quite high for -- in Cook County, Illinois, it's been quite high over time, and it peaks during the peak of the financial -- the real estate bubble. We're talking about the years 2005, 2006. The end of 2006 was the peak of the bubble in Chicago.

If we look at the percentage of people who paid the full listing price or above, you see that it's actually -- you know, it doesn't coincide very much in terms of time with the time of the credit supply. So, we see that it really shot up in '99, 2000, and then stayed -- you know, the percentage of people who paid full listing price stayed more or less constant for a number of years and then declined.

When we look at the interaction between the two, it will be an interaction, so you will have -- you know, the people who paid the full listing price and took very high leverage increased already in around 2000 and remained kind of constant until 2004, and then there is another increase. But we see already here that it's not necessarily the high supply of leverage that is necessarily linked with the likelihood of paying high prices.

Another interesting kind of stylized fact is the average leverage percent of price over listing. So, this is actually divided by the listing price, so how much of the listing price you pay. If you're at 100, it means that you paid the full listing price. If you're below, you're paying less than the listing price.

And what you can see is that the average leverage -- as expected, in a way, by this little story

I told you about I'm searching for a house and it's a budget of $100,000 and then I ended up buying at $105,000 and I need to pay this extra $5,000 five from debt -- you see that as I pay more of the listing price, I increase my leverage by about 1 percent, but there is a big discontinuity here at around the 100 percent. So, the guys who pay 99 percent of the listing price had significantly less leverage than the guys who paid the full listing price.

In black, you see the percentage of people who take 100 percent mortgage, and you can see, again, that there is a big discontinuity here. We can look also at the differences between each two buckets, just -- I guess it's the same data, same chart, just differences. What you see here, the black chart is the -- the black line is the same, but these are just differences. So, you see that between 99 and 100 percent, there is a jump in leverage of about 4 percent.

So, you know, one story could be is that these guys do not overpay. It could be -- you know, you could say, "Well, I'm a shrewd -- a shrewd buyer. I found a bargain. I'm willing to pay the full listing price. And because I know it's a bargain, it's underpriced, I'm just, you know, taking 100 percent finance, and a year from now, it's going to be worth 105 percent of what I

paid." So, this is something that I can test.

What I do is I look at repeat sale transactions, and I ask, okay, let's look at the -- you know, today's transaction and the next transaction. What happened to the prices conditional on me paying today the full listing price and leveraging it at a high leverage? Is it expected to increase or not? So, this is what we have here.

So, I can look at two sets of pairs of transactions. I can look at today's transaction relative to the previous transaction on the same property or today's transaction relative to the next transaction of the property. And what we see here is that when I compared the current to the past transaction, prices are actually increasing when I'm highly leveraged and paying the full listing price, meaning that I'm overpaying today, right? I'm overpaying -- you know, believe me that they are all -- they control for the market in the background, but what we see here is that I'm overpaying by about 3 percent relative to somebody who is not paying the full listing price and taking high leverage.

Also, if I'm comparing a future transaction to today's transaction and today I'm paying the full listing price and taking high leverage, the future

transaction is going to have a lower price by about 3 to 4 percent. So, overall, it seems that these buyers are overpaying by around 3 to 4 percent, which translate in this sample to an overpayment of, you know, around five or six thousand dollars.

Another interesting feature here is the high foreclosure rate. I don't have default, I have foreclosure in this data, but what you can see is that buyers who were highly leveraged and paid the full listing price are almost two times more likely to default on their debt relative to the guys who are only high leveraged. And you can see for the different periods, it's stronger for the earlier periods than later on, but it could be that also there are some differences in my ability to track foreclosures.

Now, you know, one response to this could be that, you know, this is perhaps accounted for in the interest that they pay. These are higher-risk borrowers. Perhaps they pay higher interest. Well, they're not. They are paying, you know, slightly lower interest than otherwise.

All right. Let me show some potential explanations. As I mentioned, I cannot test all potential stories, but I think I can point to some suspects.

All right. So, the first one is the sophistication and, you know, proxied here by income, education, perhaps financial constraints. So, we're kind of back to the original set-up. On the left-hand side, we have an indicator of whether you paid or not -- I'm sorry, whether you paid the full listing price, and then on the right-hand side, we have an interaction of a high leverage with your income, your average years of education in the zip code, and your price to income, price to income being, you know, how much leveraged you are relative to your income.

What you can see is that, you know, low-income population is more likely to pay the full listing price when they are highly leveraged. People who are less educated are likely to, again, pay the full listing price when they are highly leveraged. And people who are more financially constrained, a higher debt to income or price to income, are more likely to engage in this behavior. When we do a horse race between the three, it seems that the average number of years of education actually wins this little competition. So, this seems like some evidence for financial sophistication on behalf of buyers.

Now, I have pretty good data about real estate agents. So, I know who they are, and I can track their

historical purchases. So, what I can do here is actually look at the history of each real estate agent and see whether his history of transactions in the previous year could tell us something about the current transaction.

So, what I do, I track, for each real estate agent, the number or the fraction of transactions in which they engaged in similar behavior or their clients engaged in similar behavior, meaning being highly leveraged and paying the full listing price. And what I see that there is a very strong effect. So, in the first column, I include only the average years of education, just as a benchmark. This was from the previous regression. So, we can see it explains well as before.

But once I control for the mortgage broker and for the properties of the real estate agent, the average years of education loses its significance, and the -- whether there is an indicator of -- whether there is a mortgage broker in the transaction versus just going to a regular bank and the history of the real estate agent take over.

Another potential story is optimism. Optimism measures that, you know, if I don't buy this house today, prices are going to go up tomorrow, and here, the

question is whether this behavior is more pronounced in areas -- it's called text codes -- I'm sorry, here in this regression, I have zip codes. So, whether in zip codes in which a price -- prices were -- showed higher growth in the past year are more likely to engage in this behavior.

And what we see is that for the overall sample, the answer is yes, you are more likely to engage in this behavior, paying the full listing price and taking high leverage, in areas where price growth was very high in the previous year, and you see that it's mostly concentrated during the peak of the bubble.

So, just to conclude, the main result here is strong correlation between the propensity to pay the full listing price and high leverage. It seems that it could be explained, at least some of it, by the behavior of real estate agents, by the existence of a mortgage broker in the transaction, and so there is perhaps some sophistication and optimism.

Now, as a concluding remark, you know, you may wonder, how does the real estate agent affect the likelihood of the buyer to borrow a lot and then to pay the full listing price? So, I think a potential story is that the real estate agent pushes the buyer to pay the full listing price, because this is the way to

complete the transaction and get his fees. The buyer may not have the money to do that. So, the real estate agent many times in this particular market can help the borrower to arrange a finance.

So, real estate agents and mortgage brokers are typically networked in some way. Sometimes they are even in the same office. So, the real estate agent can tell the buyer, you know, "Don't worry about leverage or about getting the funding; I will help you to arrange the 100 percent finance. You will not even feel it."

So, that's it. Thank you.

(Applause.)

DR. NEVO: Our discussant is Karen Pence from the Board of Governors.

DR. PENCE: Thanks.

I'm really pleased to be here to discuss this paper. It's an area I've been thinking a lot about, and I think it's really, really important.

Before I start in, I do want to emphasize my disclaimer. These are my views. It's not the Federal Reserve's views. It's probably not the views of my colleagues in the audience. So, just to make that abundantly clear at the beginning.

So, I think of this paper in the context of a question I think about a lot, right? I've been at the

Federal Reserve for ten years. I've been in the mortgage area for ten years. I've had a front-row seat to a lot of excitement. And so I think a lot, like, where did we go wrong? Like, how did we end up in this extraordinary mess? And so probably the phrase that we all regret quite a bit is "The subprime is contained" sentence. It's, like, where did that go wrong, right?

And so this is my answer today. This may not be my answer a month or two from now, a year from now, it probably wasn't my answer a couple years ago, but -- so, this is a point-in-time snapshot of my thinking on this question:

I think we didn't understand what the financial system did with the mortgages after they were originated. You know, we didn't understand the extent of the slicing, the dicing, the securitization of securitizations, the major conceptual problems with the rating models. I think we didn't fully understand how integrated those securities became into the financial system in a variety of pervasive and hard-to-track ways.

The other thing is I feel like we didn't understand the externalities associated with the marginal home buyer. I think it was pretty clear, even in real time, that the subprime ARMs, per se, were not a great thing. I think what we didn't understand is the

extent to which it would go beyond the subprime borrowers.

You know, you hear a lot about problems with the products of subprime mortgages, the interest rate going up, the prepayment penalty, what have you. I think we're still sticking to the position that that wasn't the fundamental problem. So, I'm not saying these were good products; I'm not saying there weren't serious social justice concerns and equity concerns, like why are people ending up in these things. So, I'm in no way minimizing that part of the problem.

But I think the real problem was not the interest rate; it was the loan size. It was the fact that with this loan, you could -- and I put "allowed" in quotes -- you could rationalize, through your underwriting, extending a larger amount of money to a population that hadn't had it before. So, when you look at these people, it's not the interest rate. It's that they have an enormous mortgage. That's the main problem.

As Zahi kind of alludes to in this paper, these borrowers bid up the price of housing to unsustainable levels, and by doing so, imposed costs on all households. And let me just be clear, I'm not blaming the borrowers. Like, if you say this kind of thing to

our consumer affairs people, they go ballistic. So, I'm in no way blaming or saying this is the fault of the borrowers, but it is true that today, the people really suffering, it's not so much the subprime ARM borrowers, right? There's two times as many prime, fixed-rate mortgages in foreclosure as subprime ARM mortgages. I mean, this is a problem at this point way beyond the original subprime borrowers.

And just as this access to leverage kind of goosed house prices on the way up, the complete withdrawal of it from the market now is also a big problem on the downside.

So, this is a slide we look at some at the Fed. It's the credit score on prime mortgages. So, these are, you know, GSE guaranteed mortgages, and it's the percentiles of the distribution.

It starts in 2003, and you can see, the 90th percentile has stayed high, at 800. What's been extraordinary is the 10th percentile, that there's this whole extent of the credit distribution that now, for practical purposes, had credit before and now is not. So, the withdrawal of credit, the sudden collapse there has been a big problem as well. So, it's not just on the way up.

So, what I really love about this paper is it

just homes in on the marginal borrower. As Zahi said, a lot of the stuff is kind of aggregates. You kind of discern it. We had a hard time isolating it, like, who is this marginal borrower? What do they look like? What is it they're doing? What's going on? And so we've done some really fantastic data work to kind of hone in on who this borrower is.

The borrower has less education, although if I was going to pick a nit, it is a Census tract level variable, so it's a little hard to know if, like, a zip code or Census tract level education, if you're picking up education or something, correlated, but I think, as a broad point, that's right; has less financial sophistication; and there's a role for -- Zahi didn't say this, but I'll say it -- bad actor intermediaries, right? Like people -- I think it's just a fact that some people have better access to good information than other people do. And this marginal borrower overpays for the home and is more likely to default on the mortgage.

So, there is, like, kind of an externality from this borrower that I think that we haven't appreciated, and, again, I'm not, for my consumer affairs colleagues, saying that people should not have credit, that credit is a bad thing, but there are some serious externalities

that we need to recognize and figure out what to do about.

And so what do you do? And some of these issues were raised in the earlier panel. Educating the borrower is always a heroic endeavor. I've been a little scarred on this ever since I actually went to a focus group. So, the Fed at one point had a focus group. We got to watch through a one-way mirror some people that had had -- ARM borrowers, and at one point, the focus group person pulled out this thing known as the CHARM booklet, which is the Consumer Handbook for Adjustable Rate Mortgages under Regulation Z, I think. You know, you have to give it to the borrower when they apply.

They all looked at this. They're, like, "Wow, we've never seen this before. This would have been so helpful to us when we had -- when we were taking out our application." So, that was just -- I'll just say it was personally a very depressing moment, because obviously if the lenders were following the law, they did have access to it.

And I think the larger problem is no one makes money from telling a household not to buy a house, right? Like all of you, if you've bought a house and you've looked at the list of people who get paid off

before you ever get your key to the house, I mean, it's extraordinary. And if you ever write a rulemaking trying to take some of the rents away from those people, you will hear from every single one of them in the United States of America. And so there's just not a constituency of people with an incentive to say, "Don't buy this house." And I think that asymmetry is a very important factor.

I think Zahi, in the paper, let the lenders off the hook a little bit too much, right? I mean, there's the problem that the borrowers took out the loans and the problem that the lenders thought it was okay to make these loans in the first place. I think there's efforts right now to make the pain of default more salient for the lender. So, just as the borrower said, "Oh, you know, I don't have any money down, I can walk away," you could be a poorly capitalized mortgage originator, make a lot of very bad loans, and just declare bankruptcy.

So, there's efforts under way to say, "No, this default is going to be more painful for you. You can't just walk away. You are going to have to hold capital. You are going to have to retain part of the risk." Those are also very controversial.

And just my empirical observation, there's a lot of behavioral biases on the part of financial

institutions, also. There's nothing like making money to make an optimism bias become more and more pronounced.

And so the final thing, you could say, well -- government paternalism. You could say the borrower, the lender, they have optimism biases that there is just no way can ever be surmounted, and, in fact, there are countries in Asia that have LTV maximums, and there's some studies by the IMF that say these maximums in South Korea and Hong Kong have held down house price appreciation, they have reduced default, but the problem is you're then taking away credit from this marginal borrower.

And that's just something I've come to understand, as an economist, that there is a strong visceral belief that people should have a house, and you could wish that was not so, but it's there, it's a reality, it's very hard to make it go away. And so that kind of paternalism is just very, very hard to translate to the United States within our current political system.

So, those are my thoughts. I hope they were not too depressing.

(Applause.)

DR. NEVO: I think we are running a bit behind.

I'd like -- if there is one or two questions?

AUDIENCE SPEAKER: (Off mic.) I'm wondering whether there was any relationship between time on the market and the (inaudible) loan-to-value ratios being calculated.

DR. BEN-DAVID: Yeah. So, this might be a concern, right, if perhaps there is some price war between two buyers. So, there is a control there in the transaction controls, and if I exclude, also, these -- say the 14 first days, the results are still there. And, actually, these transactions actually stay longer on the market.

If you look at transactions that are bought by people who are highly leveraged and pay the full listing price, they actually take longer to complete or to close -- not to com -- to contract on.

Yes?

AUDIENCE SPEAKER: So, you have a lot of information, you say, about the real estate agents. Can you look at whether dual agency makes a difference to the frequency of this pattern?

DR. BEN-DAVID: No. I haven't looked. What would be your hypothesis?

AUDIENCE SPEAKER: (Off mic.) Well, I think the hypothesis is it would make (inaudible) more likely.

DR. BEN-DAVID: I can -- I can definitely test it. It is easy to test.

AUDIENCE SPEAKER: I saw in one of your regressions you had -- using the mortgage broker in there.

DR. BEN-DAVID: Right.

AUDIENCE SPEAKER: And I was wondering if you noticed any or looked into any systematic differences between how people use mortgage brokers versus walking in the door to the bank.

DR. BEN-DAVID: Well, this is this indicator. This is an indicator of whether there was a mortgage broker in the transaction, as opposed to a retail bank.

AUDIENCE SPEAKER: Yeah. So, the question is -- my question would be, are people who are highly leveraged, are they more likely to use a mortgage broker or are they less likely or more likely to use a bank? Does it make a difference?

DR. BEN-DAVID: Well, what I was showing there is that if they are more likely to pay the full listing price and take high leverage when they use a mortgage broker. My guess is that also the main effect, more highly -- more likely to be highly leveraged, is true.

AUDIENCE SPEAKER: (Off mic.) (Inaudible).

DR. BEN-DAVID: That's tough, right? Are you my referee?

AUDIENCE SPEAKER: (Off mic). (Inaudible).

(Laughter.)

DR. NEVO: Maybe we should --

DR. BEN-DAVID: Thank you.

DR. NEVO: So, our next speaker is Sean, who's going to talk about adverse and maybe not so adverse selection in the mortgage-backed security markets.

DR. CHU: So, thank you for including my paper in this conference. It's great to be here. This talk is about commercial mortgage-backed securities and the bundled loans that underlie them, and I want to make the standard disclaimer that the views are my own and not those of the Federal Reserve.

So, just to give you a little background, commercial mortgages are considered by a lot of people to be a second wave of the financial crisis. A lot of these loans are securitized as CMBS, and delinquencies have risen to unprecedented levels, about 9 1/2 percent for securitized loans. So, a lot of observers have blamed these bad outcomes on distorted incentives associated with securitization, both at the time of loan origination, when the lending occurs, and when the deals are bundled together as CMBS, which I call underwriting.

So, there are a lot of different areas where there may be perverse incentives, but I focus on a particular one; namely, it related to a feature of the industry, which is that a lot of the CMBS underwriters also originate loans. So, this gives them the decision whether to securitize these loans that they originated within their own deals or in deals that their competitors are doing. This gives rise to opportunities for adverse selection, and the question is, how empirically important is this?

So, here's a really oversimplified picture of how CMBS works. Most of the underwriters of CMBS are the bond underwriting shops of large investment banks and commercial banks, and when they do a deal, what they do is they buy a pool of mortgages where some of the loans come from originations that they did themselves, but they also buy originations from other lenders.

The stream of payments on principal and interest from the loans is then tranched, meaning divided into different securities with different orders of -- in order of seniority. So, when the payments start coming in, the most senior pieces get paid off first, and conversely, when the loans start to default, the most junior pieces take the first hit.

Now, most of these tranches are sold to outside

investors, but the most junior piece, which is also called the B piece, depending on the deal, goes to different types of investors, and in some cases, the underwriter retains this piece.

So, just to give you some more idea about who the market participants are and what's in the data, between 2000 and 2007, there were more than 500 CMBS deals. Just to be clear, I'm talking about the private-label deals. There are also agency deals that I won't be dealing with. And each of these deals contained, on average, about -- a bundle of, on average, 125 loans.

So, for each deal in the data, I observed what loans were in the pool, I see the tranche structure of the securities, I observe the date of the deal, and I also know who the lead underwriter is. And I put an S in parentheses, because in some cases, there were multiple lead underwriters.

I also have loan data for the loans in the deals, and for each loan, I know the characteristics of the loan origination. So, this includes variables like the loan-to-value ratio and the debt-to-income ratio and stuff like that. I know the identity of the originator, and I see the payment history on the loan through July 2010.

So, just to give you an idea, about 82 percent of the loans originated by lead underwriters end up being securitized in-house, which is what I used to refer to loans in deals that the underwriter does themselves. And also, there were a large number of stand-alone originators that only originate loans but don't do CMBS deals. That's also going to be part of the market.

So, something that I don't observe is loans that are originated but not securitized; in other words, these are loans that the lender keeps on their balance sheet.

So, here's an important stylized fact. When you look at the data, loans in CMBS deals that are originated by the underwriter -- that is, the in-house loans -- are less likely to default. So, to give you a specific number, the hazard of default is about 9 percent lower for in-house loans, controlling for observable characteristics. Also, the better performance of in-house loans mainly arises in deals containing a large share of in-house loans.

So, obviously, there are a lot of different potential drivers behind this in-house effect. To begin with, there may be nonrandom selection. So, there are basically two different margins that you want to be

thinking about. The first is the margin between loans that end up being securitized in-house versus not in-house. So, the story is the in-house loans may be chosen for particular reasons that somehow make them better.

To begin with, the underwriter may have private information about loan quality, which generates adverse selection, but also, the better performance of in-house loans may be compensating investors for a greater degree of correlation within the set of in-house loans. So, this is like a risk-return story, and basically, what it tells us is that there may be selection on unobservables even when there's no private information; in other words, all the market participants may have the same information sets, and this will still arise.

The second margin is, what's being securitized? So, the lender has to make a decision whether to keep something on balance sheet or to sell it off or they may just be, like, generating better quality or poorer quality loans at different points in time. So, for example, nonrandom selection would arise if somehow the demand for loans by competing deals is correlated with the overall quality of loans that the originator is either securitizing or keeping on balance sheet.

So, for example, if there's a shift in demand

from the competing deals and this somehow leads to a change in the proportion of loans being securitized, this would affect the degree of adverse selection.

Second, there may be a causal effect where the idea is perhaps the underwriters and the originators are exerting more effort to ensure the performance of in-house loans. So, this would be a causal story.

Obviously, disentangling these various effects is going to require some kind of a model of how the deals are put together. What I'm going to be able to distinguish between is highlighted by these two boxes. So, basically, I'm going to be distinguishing between selection at the margin between in-house versus not in-house and all of these other stories.

And as a form of shorthand at some points in the talk, I'm just going to call the blue box selection, but just keep in mind that embedded in the latter category, there may also be some kind of a selection story, only it's along a different margin. So, just keep that in mind.

And then, finally, the relative importance of these two different explanations may have potential policy ramifications. So, for example, there's currently a proposal to make the originators keep a larger slice of the CMBS securities. So, you might

think that the relative importance of the two different types of explanations would affect the effects of this policy proposal.

So, the analysis has two parts. The first is reduced form, and basically, I'm just looking at the empirical distribution of default times for loans, accounting for controls. So, this is just a simple censored hazard model, but the one twist I add is that I allow for unobserved heterogeneity in the hazard where the distribution, which is like a nonparametric distribution, depends on whether the loan is in-house. So, the difference in the distributions for in-house versus not in-house captures the effect of in-house.

And the reason why we need to account for unobserved heterogeneity in this model obviously is because later in the structural model, we need unobserved heterogeneity in order for adverse selection even to be a possibility.

Also, ultimately, I'm going to want to model the portfolio returns. So, I also need to estimate the joint distribution of loan default, and I did so using a copula. So, there is actually a fair bit of machinery that goes into the identification arguments and the estimations, but I'll just let you read about that in the paper.

I'll just quickly give you an idea of what the reduced form estimates look like. So, recalling from the previous slide, the effect of in-house is a random distribution, so I'll just give you the mean. The mean effect of this random unobserved heterogeneity for the hazard ratio on in-house loans is a factor of 0.95. So, it lowers the probability of delinquency.

Also, most of the hazard ratios for the control variables have the effects that you would expect. So, you can look at things like the loan-to-value ratio, the effect of rental income relative to monthly payments, or the occupancy rate.

I won't talk much about the joint distribution except to say that there's a fair degree of correlation both within geographic regions and property types.

So, the basic goal of the structural model, which is the second part, is to model the matching of loans, which are indexed by j, to deals, which are indexed by i. So, the key decision variable for a firm, i, is the portfolio, which I call Ji.

Now, in the data, a number of the underwriting firms actually do multiple deals. So, for tractability, I basically assumed that the underwriters are maximizing profits statically for each deal. So, sometimes I'll also just call these deals firms.

So, the key determinants of the gross profits from a particular deal are related to what's in the portfolio. So, there are a number of determinants. The first thing is we need some kind of a return distribution. So, basically what this is, is I'm taking the -- backing up a little bit. So, the return distribution is going to be implied by the default times.

And as far as what determines the default times, well, first, there are the exogenous variables, which I call wj, where j is for the loan -- the index for the loan. So, the effects of the exogenous variables I'm just going to take straight from the reduced form model, but we're also going to care about the effect of in-house status for the loans in the portfolio, which is obviously endogenously determined, and I'm going to parameterize the nonselection effect of in-house loans by a parameter, alpha naught.

We also need some kind of a rule for how the stream of payments is tranched, which I take to be an exogenous function of the return distribution. I also include a specification of demand for the bonds, and I also take that to be exogenous.

And then, finally, for each loan in the portfolios, the underwriter has a private signal, Zij,

about the quality of that loan. So, as you may have guessed already, the dependence of the portfolios, Ji, on these private signals is what drives the adverse selection.

So, basically, each of the underwriters is going to choose a portfolio of loans, Ji, from some feasible set of potential portfolios, which I define according to a set of potential trades with competing deals, i prime, and what the underwriting is going to care about is its net profits, which are equal to the gross profits that I discussed in the previous slide, plus or minus whatever transfer payments it makes for loans that are either sold or bought from other firms.

The transfer payment for a particular loan, j, between firms i and i prime, are just going to be some function of observables, which I call f(w)j, plus an unobserved error, zetaii prime j. So, basically this unobserved error is going to depend both upon the identity of the two firms that are transacting, as well as the identity of the loan.

The set of feasible trades is defined in a pretty straightforward way. So, I make the assumption that the timing of loan origination and of the deals is exogenous, and I allow for loans to be potentially matched to any deals that occur within some window of

time following the origination date, and that's how I define the feasible sets.

So, the key parameter of interest is going to be the nonselection effect of in-house, alpha naught, and once we have that, we can back out the selection effect simply by netting alpha naught from the reduced form distribution of hazards.

To give you some intuition for how alpha naught is identified, what we need is some source of exogenous variation and the propensity of loans that are originated by a particular underwriter to go into its own deal versus other deals. So, the way that model generates this variation is through two ways.

First, through variation in the set of feasible trading partners. So, if you think about it, the more potential trading partners you have, then almost mechanically, the greater the probability of that loan going into some deal that's not in-house.

Second, all of these firms have a diversification incentive. In other words, you want your portfolio to have as -- everything else equal, you want to reduce the volatility of the returns on the portfolio. So, if a particular loan has returns that are -- has characteristics that are negatively correlated with the characteristics of loans that are

being originated by your competitors, then your competitors are going to have a stronger incentive to buy your loans from you. So, that's the second source of exogenous variation.

So, to estimate this model, I don't solve for the equilibrium, but, rather, I'm going to be exploiting a set of necessary conditions based on the changes in profits based on taking the observed portfolios and then perturbing them by either adding or subtracting a loan, by having that underwriter sell or buy that loan from one of its competitors.

So, for example, if you look at the first equation, this shows the change in profits if a firm, i, sold the loan, j, to firm i prime, and basically this is going to comprise two parts, the observed change in profits, which I capture through the function r, r(Ji), and r(Ji) minus J plus an unobserved component, which is going to have both this private signal, Zij, as well as the unobserved error for the transaction payment. Similarly, we can write an equation for the change in profits if they added a loan from -- that they bought from one of the competitors.

So, you can see immediately that there's an endogeneity problem, because conditional on a loan being included in the portfolio, the expectation of this

unobservable is not going to be zero. So, as a way around this problem, Pakes, Porter, Ho, and Ishii make the observation that in a lot of cases, we can actually find linear combinations of these necessary conditions across firms and across choice alternatives such that we can either get rid of this selection problem or such that the unobservables totally drop out.

So, the basic identifying assumption I make is that each underwriter, i, has homogenous beliefs about all loans from a given originator. So, for example, if we call the originator of a loan j=k(j), then what this is saying is that i has the same private information about all loans that that originator, k, is originating. Similarly, I make a similar assumption about the unobserved errors in the transaction payments for the loans.

So, if you're familiar with Pakes, Porter, Ho, and Ishii, what this is going to allow me to do is to treat the set of loans from a given originator in the firm's portfolio as being somewhat loosely analogous to being an ordered choice.

This shows the approach a little bit more graphically. So, each of these columns is the portfolio, so -- the portfolio for firm i and firm i prime, respectively, and -- actually, I am going to skip

over this, because I want to get to the results. I already kind of told you about the estimation approach.

Oh, one more thing. So, we also can exploit moment conditions based on the total gains to trade for firms i and i prime, and the only thing I'll say about that is here we can exploit the symmetry of the transfer payments. In other words, because what i is paying to i prime is equal to what i prime receives from i, the unobservable payment drops out.

So, here are the estimates. The most important number to look at is the effect of the in-house effect, which is negative. So, this implies a hazard ratio of about 0.46. So, in other words, the nonselection effect is actually more than accounting for what we see in the reduced form.

So, what this is telling us is that selection is actually leading to in-house loans being worse, and a lot of this is being driven by -- in the model by this diversification incentive. So, the idea is if you want to include a little bit of your own loans and somehow that adds diversification benefits to the portfolio, then as the underwriter, you might be willing to tolerate slightly lower expected returns on those loans.

So, I just want to wrap up. Incentive distortions in securitization markets are a major

concern, but it's hard to quantify selection effects without imposing some kind of structure. But we don't have to do everything structurally. I am able to estimate most of the parameters directly from the data in the first stage, and the estimation can be done using moment inequalities. We don't necessarily have to solve for the full equilibrium.

And then, finally, I find some evidence that does not support the idea that the better performance of in-house loans is due to selection at the margin between in-house versus non-in-house.

(Applause.)

DR. NEVO: Thank you.

Our discussant is Ron Borzekowski. I hope I didn't butcher that too badly.

DR. BORZEKOWSKI: Beautiful. Borzekowski is perfect. You have seen this graphic before.

Thank you all for inviting me today. As we've been building the CFPB, the chance to come to a conference and indulge my academic side is actually just, you know, consumption on my part.

So, let me start with the same disclaimers as everybody else. Everything I say here today is not necessarily the view of anybody at the CFPB, nor Treasury, nor anyone that works there.

All right. Very quickly, here's the commercial fact that Sean's -- the stylized empirical fact that Sean's going to start with, commercial real estate loans and securitizations underwritten by the loan originator -- so I originate the loan and then I securitize it -- are going to perform better, ceterus paribus, than loans sold -- so I sell it to Aviv, and then Aviv securitizes it by a third party, right? And you naturally think that this is -- you know, your gut instinct should be, of course, there's an adverse selection problem.

Much like Karen, I didn't have the front-row seat to the crisis while I was at the Federal Reserve Board for ten years; I sort of had the very back-row seat, but a lot of the talk back then, you know, the first thing people would say is, how can you securitize -- you know, name your favorite thing, right? Somebody would say, of course, you're going to sell the bad stuff, and all the economists sort of thought, no, that's not right.

Markets -- you know, we sell used cars all the time. We solved this problem a long time ago. The market has mechanisms like rating agencies and like due diligence firms and like repeated gains and repeated actions and reputation effects, so that it's -- you

know, adverse selection is not -- you know, it doesn't have to happen.

But this was the same debate that was going on in the halls. I mean, a lot of people just walk up and say, "Well, this has to be happening." And you just say, "No, the markets actually handle this."

So, the economists sort of understood that this -- you know, for better or for worse, the better is, you know, we sort of think we have a rich view of the world; the worse is that no one could believe the bad stuff was getting out there as quickly as it was, at least when you got to '05 and '06, because, of course, the market would not buy lots of that stuff. Why would you ever do that?

So, in this particular paper, Sean's goal is going to be to disentangle these different effects. He wants to look at different -- you know, ex ante loan quality, some kind of adverse selection, and some ex post action by the originator. And the one note here -- I mean, he did a very good job in the talk about this, right? -- the margin here is everything is going to be securitized. The question is, do I securitize it myself or do I sell it to you to securitize, right? So, the actual decision to keep things in-house is outside the scope of this.

I am going to not, in seven minutes, try to summarize what is a lot of machinery and quite well done. So, I am going to put it -- let me first just put a broader scope on where this sits in the literature a little bit or in the crisis, because I think a lot of people in this room, it's very natural to think, okay, this is an important question. Economists always think about separating hidden action from hidden information or ex post versus ex ante concerns.

There's also an important policy area in this world. The top chart I stole from a Federal Reserve report. I think ABS Alert is the underlying data here. Don't worry about the magnitudes. This is, by year, the volume of securitizations. In green are residential mortgage-backed securities, and in red are commercial mortgage-backed securities. So, like most charts, when we're riding a financial crisis, things go up very steeply, and then they collapse. This was a running joke at the FCIC. Every time I would put a chart together, somebody would say, "Oh, great, another chart that sort of goes like this." It just goes up and it just falls off, so...

But in this case, the interesting thing, I think, is -- so, we all know about mortgages peaking in '05 and '06 and then coming down, and when Sean was

talking about a second wave of the crisis, right, commercial mortgage-backed security issuance keeps rising through '07 and then just falls off, right? So, this was -- as lots of money was coming into commercial -- into capital markets and flowing to various asset classes, this one, you know, kept going.

And the pattern in commercial mortgage-backed securities is very similar to things we saw elsewhere. So, there were more deals as the crisis went on, they were bigger deals, they were more complex deals, the covenants weakened, there was less subordination, which meant, you know, more AAA stuff supported by less.

So, all the patterns you saw everywhere else happened here as well, which make it sort of an interesting laboratory to think about the generic issues of the crisis and how we got here and this connection between the secondary markets and the primary markets, including, at the very, very end, in '07, you even saw about $40 billion in CMBS CDOs.

Karen mentioned this earlier. These are resecuritizations. So, you take some of the tranches of these commercial mortgage-backed securities that Sean showed you, you take a bunch of those particular bonds, you repackage those, rerank them, create some AAA, create some other stuff, and sell those off, even though

it sort of crept into the market in 2007, later than they had been earlier for other markets.

And, of course, then you also -- the aftermath is also true, right? Bear Stearns, when it finally went under, had a tremendous amount of Hilton debt, and that became part of Maiden Lane I, so all of you indirectly own, and Lehman and the Archstone deal as well. So, the CMBS market is intricately linked with the crisis, and a lot of the same, I think, dynamics that we saw elsewhere are relevant here. So, it's a great place to put one's efforts.

So, let me take another few lessons from the crisis, too. I think as Sean rewrites this -- and I know that papers can go through another expositional run -- I am going to throw some things at him to help think through what to put into the text. I think the model, overall, is -- it's simplified in many ways. I mean, there's some abstractions, you have to in any structural model, but it's not simple. There are a lot of moving parts, and I think he's captured exactly the right ones to help capture the features that I'm about to lay out, okay?

So, the first is the nature of the firm, and so we're going to group -- what he's going to do is he's going to treat CitiGroup -- if CitiGroup originates a

commercial mortgage and then securitizes it as CitiGroup, that that's vertically integrated. And, you know, at least part of the text talks about this, is if the same person that is making the loan is securitizing the loan and, therefore, there is no loss of information and if there is an adverse selection problem, they know it. I'm not sure -- I'd like to see more discussion.

I don't know the facts of this part of the shop, but, for example, the reason that it says CitiGroup up here, the group at Citi that would securitize mortgages almost never see subprime mortgages, never securitize Citi's subprime mortgages. In fact, there's one group at Citi that was basically buying everybody else's mortgages and securitizing them. These were residentials, right?

And there was another group that was the CDO shop, and they would not even talk to each other, and there's evidence that one side actually realized there were problems in the housing market and actually pulled back and the other kept going hog wild. So, even within CitiGroup, the information did not flow from one side of the shop to the other.

So, I can imagine in this case, for example, I think the relevant part is, if it is true that the guys that are making the real estate loans might be playing a

repeated game with the guys that are securitizing and that are facing the capital markets, and that's the mechanism why either better or worse loans are going to be there, all right? But it's not a -- this idea of a monolithic city or securitizer I would like to wrestle with a little bit.

I'm not going to -- I am going to skip, in the interest of time, the drivers of defaults, except to just say there are vintage and region effects. I want to think a little bit more about what else should be in there, whether they are adequate, and if they are not adequate, what's flowing into the error terms.

And then the other idea is that as the market -- as we got to '06 and '07, sort of just before the collapse and when this market is really heating up, still on the upswing at that point in time, the markets changed, and the markets changed in a number of ways. So, for example, on the residential side, spreads narrowed, in particular, on lower tranches, and lots of stuff could not be sold.

So, you would do a deal, you would get the fees, you could sell off a lot of the AAA, maybe 80 percent of the AAA. You wanted to sell the BBB tranches, and you couldn't do it.

Merrill had this problem left and right on their

CDO side, in particular, okay, and they started making quid pro quo deals, basically saying, "I did this deal, I'm left over with this piece I don't want to hold, you take it. Oh, by the way, if you don't take it now, you won't get my next deal," right? So, they started finding all kinds of ways to slough stuff off. In the end, they couldn't slough all of it. It was on their balance sheet when the crisis came, okay?

So, a lot of what you see in the very late years, the people -- you know, that might be securitized in-house is stuff they just could not sell to anybody else. The prices got very distorted, especially across the different parts of these capital structures.

The holders of the lower-rated tranches, again, we just need to discuss this a little bit. I mean, it is true in the early years that I held the B tranche that you put up there, because that sort of signaled to the markets that these were good loans and that I held the residual risk at the bottom of this capital structure.

That was also not true in '06 and '07. There were people that were buying these in order to bet against them and who didn't care and who would buy anything, and that also distorts incentives across the AAA and the B that's supporting it. So, I just -- you

just want to prove, for example, or try to argue, at least, that in the '06 or '07 period, the same market dynamic that was happening in 2000, 2002, 2004, where you hold the B piece, is still valid, or cut the hazard short. I know you lose a lot of defaults that way, but that may be another way to handle this.

And then similarly, I'll give you a little bit of evidence about hold versus buy. Washington Mutual got into the option ARM product in 2003. It was a very conscious choice of theirs. This was when they started making option ARM residential mortgages. They were securitizing all of them. They looked at their book in '06, and they said, "Wow, these are really profitable. We should hold onto them."

We don't have WaMu anymore, largely because those -- they decided the market was not paying them enough for these things, they held them in their portfolio, and when the crisis came, those are some of the loans that actually took them down very quickly. So, the incentives in that very last period in '06 and '07, even the behaviors that your structural model is trying to capture, got distorted in a lot of ways, which is another way of saying I think you may want to shrink the hazard -- the time over which you do this.

But like I said, the CMBS market is a great

place to look. These are important issues. I think we have a lot to still learn about, you know, what came to be.

So, thank you very much.

(Applause.)

DR. NEVO: Thank you. Again, I think we're over time, but still, if there's one or two quick questions. Any quick response you wanted to discuss? Just anything quick.

DR. CHU: Nothing too specific. Those comments were really helpful. I think a lot of your comments had to do with how are the incentives changing over time, and I have, you know, a lot of time dummies in there. I think I need to think a little bit more about, you know, whether those are adequately capturing these changes in incentives. So, those are some pretty helpful comments. Thanks.

DR. NEVO: Okay. Let's move on to the last paper, which is J.F. Houde from the University of Wisconsin.

DR. HOUDE: Okay. Thank you very much for putting the paper on the program. This is joint work with Jason Allen of the Bank of Canada and Robert Clark. I should say, I thought I would be the only one not to have a disclaimer, but then I realized that (inaudible)

actually forced me to put something on it. So, yes, the Bank of Canada doesn't know what we're doing, okay?

Okay. So, quickly, this is where we're coming from with this -- with this paper. I mean, just -- this is sort of a standard observation. There is many markets where concentration is an issue that are not the standard posted price market, right? So, there's a lot of markets where prices are negotiated, where you have to haggle to get to actually a better deal, and a lot of those markets also have this kind of search feature, where, you know, not everybody is aware of all the options, not everybody is considering all the options. These are just examples. Consumers loans, we are going to look at mortgage, so that's what we're going to study today, but this is quite prevalent, where a lot of antitrust questions are relevant.

Now, the reason we're interested in those, in part, is because the standard method that we have to sort of measure market power in those markets don't really apply here. I mean, I'm sort of referring to the standard discrete choice model that we typically use. I mean, these are two reasons why these markets don't fit necessarily that framework.

While on the one hand consumers don't necessarily search all the options, so you have the

consumer choice set, if you will, and the other option -- the other problem, at least in our context, the bigger problem is that you only see the transaction price. That's the only -- you don't see the offers that people consider before purchasing, and, therefore, you sort of need a model to fill in these counterfactual prices, so that the price you see is not, you know, sort of Bertrand-Nash price that you would otherwise use, okay?

So, the objective here, in some sense, what we're kind of going to, is to develop a model that will -- may make how these markets work and estimate it, and in that case, it is going to be the mortgage markets. So, in the interest of time, I didn't cite anybody here, so there is no reference, but there's a big literature on search and negotiation, obviously, and labor and IO, and the IO mainly on the theory side, there has not been that much work.

The model we are going to present is sort of like a labor search model, where people actually bargain for their wage. So, we are going to borrow somewhat from the labor literature, although everything is going to be much more static than what you would see in the labor literature.

Okay. So, before I'm getting people confused,

this is not a standard U.S. mortgage market. Here's five reasons why. So, we have Canada. Canada is a lot more concentrated than the U.S., although with all these mergers and the buyouts, it seems like the U.S. market is converging toward that.

So, the market is dominated by eight national lenders that nowadays control about 80 percent of new mortgage, okay? So, they are mostly present in all provinces, although there is some more regional players.

The other thing that is going to be really good for us, that will simplify things, is that we're looking at a subsegment of the market, which is the main part of where new home buyers fit in, where lending is really simple to understand, okay? So, this is -- we're looking at insured mortgages.

So, in Canada, if you don't have a 25 percent down payment, you have to buy insurance from a government-backed institution, that says the rule, okay? So, they are very transparent rules. And, therefore, the lenders don't face any default risk, and we don't really see that much default anyway, because those rules are fairly tight, okay?

So, we are going to look at that subsegment of the market in which sort of lending, in some sense, is kind of in the background but not explicitly modeled.

The other thing that I've sort of started with is that this is a completely decentralized market. So, you see posted price, but nobody is really paying the posted price. There's no variation across regions in the posted price, and there is no variation across lenders in the posted price. So, the posted price are there, more or less like in the housing markets, where nobody -- where very few people are paying the posted price, and people have to negotiate to get discounts.

And it's decentralized in the sense that these national banks delegate that authority to branch managers, who have the responsibility to issue discounts and bilaterally negotiate that with consumers, okay? And as a result, you see a lot of dispersion. There's a little bit of typos. This is not 0.5 basis points; this is 50 basis points. So that if you look at, in a given week, the average standard deviation and ways that people are paying varies quite a bit.

So, 50 basis point standard deviation, if you (inaudible) range, if the posted rate is 5, on a typical week, you will see a range between 3 and 6, okay? And these are for complete homogeneous contracts. We are going to look at a very standard fixed rate, common amortization period, okay? So (inaudible) margin is good, but still tons of dispersion in those who

negotiate price.

Now, the other feature that is important is that not everybody is searching equally. So, we don't have data on -- we know the transaction price and we know the bank, but we don't know if people search or not. But there's some aggregate survey evidence suggesting that about half people get the first quote -- buy from the first lender that they visit, okay? So, they normally get one quote, okay? It doesn't mean that they don't haggle with them. It could be a good quote, but on average, half people search, half people don't search.

And the last fact is that there's a lot of loyalty. These are national banks, so these are multiproduct firms. They are offering day-to-day banking and lending, and the vast majority of consumers stick -- get a loan from the bank they have day-to-day banking activities with, okay? So, in our data, nearly 80 percent of consumers do that, and the other thing is -- well, so, there's two points on the loyalty.

Most people -- 80 percent of people get a quote from their home bank, okay, and the second fact is that overall, in the country, about 60 to 70 percent of consumers combine the two together, okay? So, a lot of loyalty.

Now, these facts kind of motivate us in asking

this question. So, as I say, the overall goal is to try to measure market power, and we're addressing here and trying to measure the market power of these national banks. So, where is the market power coming from and how big it is? And we're going to focus on two channels, okay?

We're going to label the first one incumbency advantage, and that's going to be referring to the search frictions, and second one, called differentiation, that's going to be coming from the fact that these are multiproduct firms and there might be complementarities between the services that they're offering.

So, what I mean by incumbency advantage, well, I'm going to say that big national banks that have large consumer base have advantage in the sense that they receive, essentially, a lot of consumers who visit them for the first time to get a quote. So, a lot of -- you know, if you go -- start your search process at the home bank, you call them to get a quote, and you are a high search cost consumers, the home bank, knowing that, will offer you somewhat of a bad quote initially, and this will give the advantage to banks with large consumer base to be able to essentially sell to a larger fraction of nonsearchers or nonshoppers in the way I am going to

model that.

And then the second aspect of the market power, the national bank, if you will, because they sell multiple goods, to the extent consumers value having everything together, this is going to give an advantage, which is going to create value in some sense to banks with large brand network and for essentially the home bank. The home bank will have an advantage of selling those -- selling to those consumers.

Now, where is that coming from? Well, you could think it's coming just from switching costs, that consumers really like to combine these things, and when they shop for the mortgage, they would incur a cost of switching their day-to-day banking account to the other bank; or it could just be straight complementarities, that I get a better line of credit if I have my checking account with the same bank, okay?

So, we're not going to necessarily -- we don't see much on the other side of these transactions, but in the back of your mind, this is what you should be thinking about when I talk about that.

So, the empirical -- sort of the goal of the empirical analysis will be to sort of disentangle -- so, we see a premium, people who don't switch institutions pay more, and we're going to try to disentangle where is

that coming from. Is it because of high search costs or is it because of these complementarities, okay? And the results are -- the paper is in somewhat a rough shape at this point, so open to comment, but don't expect final results necessarily, okay?

Okay. So, here's the outline. I'll try to squeeze everything into 20 minutes. We're going to talk a little bit about the data in the market, but I want to spend more time on the model, so I will probably fly through this first section and then show you a little bit what they likely look like, and then the results.

Okay. As I said, the market is fairly concentrated. This was not always the case, and during the nineties, we've seen a big wave of mergers, where the Bank of Canada revised the regulation that defines the services that a bank can offer and essentially allowed the banks to be -- become -- to enter this -- the market for trusts and the market for mortgages. They were already there, but they were not dominant, and throughout the nineties, they became dominant.

So, you saw a lot of trusts -- small trust companies and insurance companies being acquired by the banks, okay? So, nowadays, 80 percent of the new contracts are issued by the main institutions. The trusts don't have much market anymore. And there's a

few foreign banks and small credit unions that control about 10 percent of the market, okay? So, we're going to focus in this paper only on the stable period between '99 and 2004, okay?

Now, how the market works, well, as I said, there is really two segments. There is this insured market and the uninsured market? We're going to focus on the insured market, and -- well, part of the reason is because we have the data from it, but it's also because it simplifies our task. We're going to focus really on the (inaudible) shopping decision and not so much a lending decision, and we're going to assume that essentially consumers who bank with a certain bank in our data would have access to a loan at another bank because they are covered by the government insurance contract, okay? And this is by far the largest segment if we're talking about new home buyers, which is sort of the population we're interested in.

And as I said, this is a market that is decentralized, national posted price, very little dispersion, but everything is determined at the local level, and within -- so, two branch managers from the same bank are not going to compete. So, you are not going to -- if you get an offer from TD, you are not going to be able to get another TD offer that will lower

your interest rate, okay?

Where does the data come from? So, what we did is that we went directly to the two insurance companies that issue those insurance contracts. Up to '96, CMHC is the public insurance company that was the only one in the market. Genworth entered in '95. So, we got data from both companies, which is sort of a running sample of all contracts that were issued between '92 and 2004, although today I'm going to be focusing only on the four-year period.

And what's in the data, well, we know the contract type, so we know everything the insurance needs to know to price the contract. So, we know the contract terms, the financial characteristics of the lender. Through a series of confidentiality agreements, we managed to get the lender information. And we also know where the house is located. So, this is relatively crude.

This is at the FSA level, which is the first three letters of the postal code, a little bit bigger than the zip code in the U.S. So, there's about 1500 locations throughout the country, okay? You can think of it as a small city or a Census tract within a big city, okay?

And, interestingly, we know if you had a prior

relationship with the lender you're dealing with, which is how I'm going to be measuring whether you switch institutions or not, okay? And today, I'm going to be focused on a very restricted sample or somewhat restricted sample, because I want to look at very homogenous contracts. So, we're going to look at this period between '99 and 2004, and then focus on the most common contract, which is about 75 percent or 80 percent of consumers during that period were choosing that contract, which is a 25 years amortization and a five-year fixed rate.

So, there's -- in Canada, you don't see these very long fixed-rate contracts. Almost everybody buys a five-year fixed rate. You see more variable rates starting to pick up later on, but during the sample five years, the standard product. And we are only going to look at new mortgages, no re-fis.

Okay. And then there's -- a big share of the market is served by brokers, and I won't talk about them at all today. So, I'm just dropping them from the sample, okay?

This is just to give you an idea of how disperse prices are. This is not necessarily representative of all the years, because there's -- in 2000, it seems like -- over time, there's been fewer people who pay the

posted rate, but in 2000, we see about 10 percent of people who are paying the posted rate, and then the rest are paying above, a little bit below. Part of the reason is that we know when the closing date of the transaction, but we don't know when you actually negotiate the rate. So, there is measurement error in the actual posted rate, but when you see this dispersion between essentially -- this is the deviation from the posted rate, so between zero and two is where most of the data lies, okay?

These are just summary statistics. Skip that. And in the paper, there's -- yes, okay. In the paper, there's a little bit more of that. There's a couple of reduced form relationships I want you to think about before I get to the model. People who are loyal to their bank pay more, okay? So, switchers pay nearly eight basis points less than loyal consumers.

If you live in a neighborhood where you have access to many lenders, you pay less, right? So, the number of lenders in the neighborhood of the new house sort of decrease the rate you are paying. And if you bank -- if you're choosing a contract from a large network institution, you also pay more, okay? So, if you bank with the dominant player in your local area, you pay more versus the smaller guys.

And then on the right, there's -- we don't see people searching. In the model, as you'll see in a second, we're going to, you know, proxy for -- well, switching is going to be an indicator of whether you search or not, and the regression sort of suggests the thing you should expect, that people who are financing a big loan are more likely to search or switch. People who are choosing a small institution are more likely to switch.

Okay, let's get to the model. So, there's a lot of assumptions, so -- as always. So, here they are. So, what we want to do here is we are trying to measure, you know, how big these search costs are and how big these premiums are for valuing the home bank, but the challenge is that we don't observe necessarily search. We observe some aggregate measures of search, where several people switch or not, and we observe the transaction rate and the lender, okay?

So, given this constraint, this is what we have to make assumptions. So, we're going to assume that everybody is affiliated with a home bank, so that's not too hard. To define the choice set of consumers, we're going to say that people shop locally. So, we're going to talk about it afterwards. This is true in the data, that most people -- the distance from the closest branch

of the bank you choose is much smaller than anything else, and on average, people seem to buy within about two kilometers of the new house that they buy.

Okay. So, that's the first assumptions. We're going to assume that everything is done locally, and at most, you're getting quotes from the banks that are located within ten kilometers of the house that you buy, okay? So, think of this as the city where you live, and that's going to vary across consumers, okay?

But not everybody does that, so people who are searching get quotes from these banks, and people who are not searching get a quote only from one bank, and that quote from -- the first quote is going to be coming from your home bank if you have one in your neighborhood, but if you move city or something, then it's possible that your home bank is not there, and at that point, we are going to assume that you just pick randomly according to the branch distribution, okay?

Now, searching is costly. So, we are going to model it as a private value shock. So, cap i is going to be the search cost -- give me five -- no? Okay. So, cap i is the search cost, is the cost of collecting these additional offers, and that's probably of value to the consumer. And the timing is as I described. You first go to your home bank, get a quote. If you like

it, you accept it. If you don't, you pay the cost, and then you go shopping.

Now, we're going to model the shopping in a way that is going to be retractable, and we will approximate the idea that you go to bank one, you get a quote, and then you go to bank two, and you allow bank two to respond, and then you go to bank one to allow him to respond, and you do that until -- you know, until nobody wants to respond. So, this is very much like how our auction would work, and the auction actually helps us a lot, because it's a lot more tractable.

So, we are going to model this negotiation process as an ascending auction, where consumers get quotes from the banks in their neighborhood, okay? So, if you pay the search costs, you go to the auction.

Preferences. So, there's a -- because we want to talk about differentiation, people will value, to some extent, the characteristics of the bank. So, that's going to be theta, and they are going to value the contracts like this, and then banks are going to make standard profit. C is going to be reduced form, that function of the financial characteristics of the bank, and then Ui is sort of a private value shock that the banks get.

This is sort of a match value for the bank of

having that consumer in the bank as a consumer, and that might differ a little bit across vendors. Okay, the V is the surplus -- the transaction surplus.

Now, to show the model backward, at the auction stage, it is an ascending auction with differentiations of -- I mean, just skip through this. Essentially, the highest value bank or the high surplus will win, and that bank will pay the second highest valuation, the second highest surplus available in the market.

And then, going backward, in the initial stage, what the home bank will do is the home bank is a bit like a monopolist facing a random demand, right? So, it wants to retain high search cost consumers, so, therefore, will maximize its profit, where H is the probability of a search, okay? And the probability of a search is sort of incorporating the result of the auction or the expectation of the auction.

But the model -- there is two cases that I am going to be showing you results for. The first case is sort of what I call the full information model, where consumers know all the match values, know all the use, there is no uncertainty. In this case, the model becomes very simple, and the essential quote is actually linear in the cost.

In the more general case, which is very similar

but a little bit more numerically complicated, the quote is such that if you have a good match, banks will offer a good rate, up to a constant. If the home bank knows it's going to win the auction for sure, essentially it will just give you a constant and won't pass this extra benefit. And we make those -- those functions four assumptions.

So, U is going to be -- I told you I didn't cite anybody, but I do cite somebody. So, I -- we have this match value, a U, which is going to be assumed to be extreme value, and this is actually very useful in our setting, because we have heterogenous -- it's an auction with heterogenous bidders, and everything becomes much simpler.

Let's skip through this and just show you the results. So, we're getting results for the incomplete version of the model, but the standard errors are not there, because minus 2, the likelihood of seeing behavior, but anyway, the main result is that -- well, not too surprising, here we see big search costs. So, everything here should be expressed in hundred dollars. If we can actually go to the next slide, where I put dollar values.

So, we're estimating fairly big search costs. So, the average monthly payment here is about a thousand

dollars, and we're estimating that it costs a consumer -- everybody about $20 to search, but this is, you know, relative to a (inaudible) multiplied by 60 to figure out, you know, over the life of the contract, what would this represent. And there's quite of heterogeneity around this. So, not -- the average consumer is paying (inaudible) $50 per month for searching.

So, this is big. A little bit smaller than that is the home bank premium, but still quite large. So, we see a lot of loyalty. So, that loyalty is explained both by search friction and this home bank premium. And we're estimating, depending on the model, that people value -- that they would be essentially willing to take an extra $40 on their monthly payment to be able to combine everything at the same hood and -- and, okay, and that's it.

So, that's basically it. So, I mean, this is still work in progress, so we don't have final results. This is where we're going. There's a lot of, you know, improvements to be done in the model. One thing we have left out that is quite important here is the brokers and the financial intermediaries, and that's going to come.

Thank you.

(Applause.)

DR. NEVO: Okay. Our final discussant is Ken. So, we have to have someone from the Board on each paper.

DR. BREVOOT: Actually, I was going to say, I have really two disclaimers. The first is the one you've heard here all day today about how my opinions are my own and not necessarily those of the Federal Reserve, and I was also going to point out, since it has come up, that Ron, Karen, and I started within about a couple months of each other at the Fed, so I was going to tell you that in my ten years at the Fed, I sort of had a view -- I had a seat sort of in the balcony of the Fed, watching the housing crisis, and then it felt like somebody pushed me and I have been falling through the mezzanine, and now I find myself feeling like I'm lying flat on my back somewhere in the orchestra wondering what exactly just happened. So, please use that to interpret my results, as you see fit.

As J.F. just said, this paper that I'm discussing really has preliminary results, so a lot of my discussion is going to be more focused on sort of broader, high-level issues that I think the paper presents and some of the topics that I think might be of particular concern to the FTC and people at the Board and places like that.

So, I'll start out with what I actually like about this paper, and there is actually a lot to like in this paper. Particularly, I like the structural approach, because as an econometric methodology, I think it's very valuable, but more fundamentally, I think the idea of looking at shopping for mortgage products as a search process really is a necessity. I think it's really the right way to actually look at the way people go about purchasing mortgages, and I think it helps explain some of the things that go on in the mortgage market or have gone on in the mortgage market in a way that other types of models fundamentally just can't.

For example, to take you back to the heady days of 2007, where we had -- you know, the mortgage market was starting to slow down dramatically. We had the implodometer sort of helping us keep track of exactly how many subprime lenders had gone out of business. One of the things that we saw in that year when we started looking at the HMDA data was that what you found is that almost across the board, when the activity was down -- it was down I think about 20-something percent overall; amongst Hispanics it was down 75 percent, largely because Hispanics were disproportionately located in California, Arizona, Nevada, Florida, the places where the market was really being hit heavily; and it was down

over 50 percent among African-Americans.

So, at the Board, we took a really careful look at sort of trying to understand why it was that minorities were being hit by this so much more intensively than it was other portions of the population. And what we found, what was driving the African-American decline, in particular, was that it had to do a lot with credit risk, and, in particular, if you controlled for the pricing of loans, if you looked at the distribution of APRs between 2002 to 2006 and 2007, what you found is that it was really the highest end of the credit risk spectrum, the high APRs, where you'd see the market had essentially just collapsed and that African-Americans were disproportionately located there.

Now, why I bore you with all of this is the fact that what we found at the bottom end is not something that we expected. In a market where everything was collapsing, where lending activity was down substantially, we actually saw an increase of lending to African-Americans in 2007 at low interest rates; people who were getting offers just around the prime cut-off. And this shocked us tremendously, because this was really the only increase you could find in the 2007 data.

And when we looked into this, what we found is

that if you looked specifically at those neighborhoods where all these subprime lenders had disappeared during the course of 2007, the 169 or whatever it was that had disappeared, that growth was coming into this market. And the way we interpreted this was that fundamentally, what had happened is that in 2006 and earlier years, where you had these subprime lenders that were in these markets, that were doing a lot of lending activity in these neighborhoods, African-American borrowers, in particular, were going to these lenders, and when they disappeared, they had to go somewhere else, and when they went somewhere else, they got dramatically better prices.

So, it fundamentally matters which door you walk in first, and that to understand what's happening in pricing, particularly for the FTC and the Fed and others who are interested in sort of lending to disadvantaged populations, a search model, I think, is fundamentally exactly the right way to go about it. So, I was actually really excited to read this paper.

Now, the downside. When I first read the paper, I sort of had an almost visceral negative reaction, because occasionally what you will do is you will see people who estimate structural models, they will look at structural models in banking, in particular, and they

will look at the Berry-Levinsohn-Pakes methodology and things like that, and they will say, "Wow, this has only been applied to deposit markets." So, they'll write a paper saying, "All right, I'm now going to apply it to credit markets. I'll do, you know, mortgages or things like that." And they write these papers, and the papers say absolutely nothing about credit risk.

But credit risk and prepayment risk are fundamentally so essential to the pricing and consumer choices about credit that, really, I would argue such papers are really only going to make marginal contributions, and the fact that this paper didn't say a whole lot about credit risk sort of bothered me greatly. But then I realized it was Canada we were talking about.

And so I went -- as anybody would do in that situation, I went to Google, and Google actually referred me to another paper, I think the three authors that was published on the Bank of Canada, where they actually laid out exactly what the Canadian banking market for mortgages was like, which was enormously helpful, because what it did is it actually convinced me that credit risk in this market may not be as important as I sort of would have expected going in, largely because what you have in Canada -- and J.F. talked about this a little bit -- is you actually have mortgage

insurance, right? So, all of these mortgages that were made, they were insured, and unlike in the United States, where private mortgage insurance only covers portions of the loss, the insurance actually covers everything.

What's more, the costs of the insurance are paid as a lump sum up front, they're tacked onto the amount of the loan, so they don't directly factor into the interest rate, and in case you may have been worried that some of these PMI companies are going to go out of business because they will only really get their money from the mortgage market, and if the whole mortgage market falls down, what do they do, the whole industry was basically backstopped by the Canadian Government.

So, I do think that this sort of -- as I tell it here, how I learned to stop worrying and actually really liked the paper is that I do think there's reason to believe, perhaps, that in this market, the Canadian mortgage market, it may be appropriate to actually sort of decide not to really look at credit risk more dramatically.

I think the downside of this paper for people like myself, who tend to worry about mortgage markets and really are struggling to understand what's happening in the U.S. mortgage market right now, what are the

frictions, why is this not working as well as it might otherwise, I think the paper is a little bit less applicable to that than it would be otherwise, which is not to say it is not a contribution, but it probably has more in common with, say, the structural models that have been applied to auto lending and to breakfast cereals and other products like that.

Now, the reason I said like instead of love, for those of you who would catch the Dr. Strangelove reference, is that what I'm showing here is Table 6 that appears in the paper, and this is the reduced form estimate of the margin of interest rates, and one of the things that I was struck by is that if you look at the relationship between, in particular, LTV and FICO, what you tend to see is that these have relationships with the prices that we observe that are sort of in line with what we would have expected if credit risk had been important, right?

You see, the FICO is declining, so a higher FICO score gets you a lower interest rate. LTVs go up pretty much monotonically, interest rate, and the question was, why would this be the case in a market where you had no credit risk? And so this is sort of making me a little bit concerned that there may still be some sort of credit risk out there that is perhaps not being entirely

accounted for, and that would cause me concerns.

I know in some cases, for example, people used to argue that credit risk did not matter in the GSE market because, in fact, the other GSEs were the ones holding the credit risk. That is actually really not true, and I think there are reasons why credit risk still is priced in these products. But that would be one thing that I would ask the author to be a little bit careful of, and perhaps if he can convince me that, in fact, there is something else driving this, it would be pretty valuable.

So, in my final minute and a half or what I have left, I'll just throw in some additional things that I think would be really helpful to pay for my address. One is the issue of is credit risk really not that important, because I do think that when you're looking at mortgage markets, maybe not in Canada, but particularly in the U.S., if you're not talking about credit risk, you really can't understand how these things are priced or how it is that people come to choose where they go.

Also -- and this I didn't really throw up here as a comment until I went back and had to type the title of paper back into the PowerPoint slides -- is this really a differentiated product and how? One of the

things I think is really fascinating about the mortgage market is that to a certain extent, the ability of people to search across different products is going to be tied to the characteristics of the products themselves, and one of the interesting issues here that potentially he may be able to look at is sort of how do you set the characteristics of the product perhaps to boost your market power by inhibiting the searching of these products, right?

Even in the heady days of 2006 when the subprime market was really going great guns, if you went to one of the websites to try to shop for an interest rate for a mortgage loan and you had a prime credit score, you were fine. You would go to the website, you would type in your credit score. It would say, "Your interest rate is going to be this." If you put in a subprime credit score, it would say, "Give us a call."

So that if you were a subprime borrower, in particular, it was much harder to shop. They tended to give you products that were much less standardized, had lots of different features in them, all of which were sort of designed to sort of help or hinder a borrower's ability to do switching costs.

And so that sort of plays into sort of the third comment or third thing, what actually is it that

determines these switching costs, right? You could assume that there is a distribution of switching costs and what we are seeing is draws from that, but I think there's a richer story here, which says if somebody is concerned with sort of the public policy issues here and the functioning of the mortgage market, I think it would be very valuable to sort of get a better idea of exactly what it is that's driving that.

But the main take-away is I really think this is a very carefully done and very promising paper. I think it's looking at the mortgage market exactly the right way, and I really look forward to the next draft. Thank you.

DR. NEVO: Okay. We are officially to lunchtime. I mention that because I will open the floor to questions, but you are standing between 87 hungry people and lunch. So, are there any questions?

(No response.)

DR. NEVO: Great. I would just like to conclude and thank the authors and discussants for what I think was a beautiful session. Thank you.

(Applause.)

(Whereupon, at 12:43 p.m., a lunch recess was taken.)

AFTERNOON SESSION

(1:11 p.m.)

DR. O'CONNOR: Hello. My name is Jason O'Connor. I'm an economist here at the FTC, and it is my pleasure to introduce David Dranove. David is the Walter McNerney Professor of Health Industry Management at Northwestern's Kellogg School of Management. He is also the director of the school's Center for Health Center Market Economics. Professor Dranove's research focuses on problems in industrial organization and business strategy, with an emphasis on the healthcare industry.

He has published numerous papers on topics such as the relevance of hospital quality report cards, antitrust enforcement in healthcare markets, and the effects of provider affiliation and integration on patient flows and costs. He has written five books, including The Economic Evolution of American Healthcare and What's Your Life Worth? His textbook, The Economics Strategy, is used by leading business schools around the world, and his latest book, Code Red, was published by Princeton University Press in 2008.

He also shares his insights on healthcare policy issues through a blog of the same name.

(Applause.)

KEYNOTE ADDRESS

DR. DRANOVE: Thank you.

I want to thank the FTC for inviting me here this afternoon and giving me an opportunity to read papers and help organize the first session this afternoon on antitrust.

Before we get to our speakers, I am going to tell a cat and dog story. Why it's called a cat and dog story will be clear to all of you if you manage to stay awake through the course of my presentation.

Merger analysis in antitrust is often delegated to industrial organization economists, which is why industrial organization economists wear much nicer clothes than most other economists. IO provides theoretical models that form the foundation for ex post empirical studies.

Sometimes these theories are sufficient to make predictions about future mergers, but sometimes merger analysis tends to run ahead of the theory. Some analyses rely on ad hoc approaches, with shaky theoretical foundations, which often, because of a reliance on precedent -- and I'll try to show you how that's worked in healthcare antitrust -- we use existing approaches simply because they exist.

There is also a preference for simple theories

with simple empirical implications, even if the theories do not fit the facts. On more than one occasion, I have lawyers immediately shut their eyes and ignore me when I started talking in Greek letters.

So, there is a long history of IO economists trying to develop and empirically implement theories of competition that do try to match the institutions and fit the facts, and Cournot's model is the first and best-known example, I think. It's a sample game theoretical model which provides the foundation for using the Herfindahl Index to predict merger effects, and, indeed, to this day, agency guidelines cite the HHI, and they talk about the HHI ranges in which markets are competitive or not competitive.

And while it's not definitive of how the agencies view mergers, they have a large practical impact. Nearly any organization that's trying to merge, one of the first things they're going to do is try to document a small impact on the HHI.

Now, the Cournot model, while it tells us that the HHI is useful, it describes a world with homogenous goods, where firms make simultaneous capacity or output decisions, depending on how you want to view the model, and surely this is an approximation, at best, of most markets. Even so, traditional structure-conduct-

performance regression studies do suggest that HHI could be a decent predictor of pricing in many sectors, which is perhaps one reason why the Cournot model and the HHI are used so often. It's directionally correct, but that does not mean that we should foreclose further theoretical and empirical study.

Measuring the HHI is clearly, though, not enough, because if you are going to measure the HHI, you'll have to identify who the firms are in the market, and that means you have to define the market. The guidelines from the DOJ and FTC tell us that we should implement the small but significant nontransitory increase in price -- or SSNIP -- test. Before this, market definition was pretty ad hoc. We would rely on SAC codes or county boundaries, stylized evidence about substitution patterns. So, the SSNIP test now gives us a firmer guide to develop our theories, but it doesn't tell us exactly what we're supposed to do. It's just that we should do something so that we can implement this test.

It could tell us that the impact of a hypothetical increase in the market HHIs is what we should be -- I'm sorry.

Structured economist performance studies, excuse me, tell us that what we should do is look at the

increase in the HHI, okay, and so we often will do that, and you'll see that in analysis, we run price as a function of HHI, we get a coefficient, we then see what the HHI change will be, and we compute the predicted increase in price. But that's all completely circular, because we've assumed that the simple structure-conduct-performance regression was the right regression to predict merger effects in the first place, and we assume we've got all the right firms, and so we've computed the HHI correctly.

We need to do better than that, and this is where structural modeling comes in, and after I finish this slide, I am going to try to transition to tell you about how all this has worked in hospital mergers. Structural modeling offers us several advantages for merger analysis. We can tailor our assumptions about market conduct to the institutions. We don't have to take a model in which we assume that firms are choosing capacity and setting those capacity choices simultaneously and hold our nose and say, "Well, I know it stinks, but we're going to apply it somewhere else."

The model can specify in theory and recover from the data the conduct parameters that you're most interested in in order to predict merger effects. It allows us a way to avoid ad hoc market definitions. We

could, in principle, take every firm in the world and throw it into a structural model and then implement the SSNIP test. We don't have to decide beforehand how to compute the HHI and which firms to put into that calculation.

Armed with our conduct parameters, we can do analyses of hypothetical scenarios. If A merges with B, what do we predict will happen? And thus, we can do the SSNIP test.

To paraphrase my one-time classmate, Garth Saloner, structural models can provide an "audit trail" that allows us to better understand how specific assumptions generate specific conclusions, while ad hoc specifications and the FCP model essentially keep everything hidden in a black box.

But structural models do have problems. Models that exactly describe the real world would be intractable. In the full version of a model I'm going to describe in just a few minutes, you would have to solve 16 million simultaneous equations, which is not only more than we can figure out how to solve as economists, but it's probably more than the decision-makers out in the marketplace are able to solve.

The models have to be developed with an eye

towards available data. The well-known Berry-Levinsohn-Pakes model for studying demand in differentiated goods markets, for example, exists because we don't normally have transactions-level data. So, they take data at the market level, and they look at shares of firms at the market level. If we had transactions-level data, we could, in principle, do better.

As a result, models represent a compromise between describing the real world and the ease of implementation, and as a result, it can be difficult to ascertain where the results come from, and the results can be sensitive to the specific assumptions about market conduct in ways that are much easier for the other side's economist to point out and challenge in the courtroom. And explaining all of this to lawyers, judges, and juries can, therefore, be very difficult.

As a case in point, let me talk about hospital merger analysis. Nearly all of these issues have arisen in this context. This is an industry that's undergone massive consolidation, with local mergers starting in the 1980s, being commonplace in the nineties, and often, you'd see markets with three hospitals merging into two or two hospitals merging into one, and yet, nearly every time these mergers were challenged, the antitrust

agencies lost.

So, a lot of economists turn their attention to hospital markets, and as we've done so, our analytic methods have evolved from ad hoc to structure-conduct-performance to structural. In the early days, back about 1990, the highlight was the Rockport, Illinois, merger decision. Prior to then, there were few hospital mergers and little enforcement activity. There were structure-performance studies at that time, and they suggested that more competition led to higher prices, and that actually has been in the mind-set of people in this field ever since.

But in U.S. versus Rockport Memorial Hospital, the Department of Justice blocked the proposed merger, and they did so by using a tool for defining markets that was developed by Elzinga and Hogarty in a series of papers that they published over a decade earlier, papers that they showed that you could use the flows of goods and services to define the markets for commodity products.

Using a similar methodology to look at the flows of patients, the DOJ showed that Rockport and Chicago were separate markets -- they're about 90 miles apart -- and once you show that they were separate markets, you now have the merging hospitals in Rockport gaining a

substantial amount of market share. The HHI increases substantially as a result.

In the model, the market is well defined if the inflows of patients from outside the market into the market represent less than 10 to 25 percent of the patients who are treated in the local hospitals and the outflows of patients who live in that market to hospitals outside the market represent less than 10 to 25 percent of the people who live in that market. That's the Elzinga-Hogarty test.

Well, the DOJ won the battle, but they may have lost the war, because between 1994 and 2001, the DOJ and the FTC lost seven consecutive merger cases, including mergers in Joplin, Missouri, and Dubuque, Iowa, which if you looked at a map and you looked at the little pins on the map, you would say, "My God, those are mergers to monopoly."

In all but one of these cases, the Court's decision hinged on an Elzinga-Hogarty-style analysis or related critical loss analysis. That suggests that these markets were very large. For example, the hospitals in Dubuque, Iowa, were -- claimed to be competing with hospitals in Iowa City, which was 70 miles away. This came from the application of Elzinga-Hogarty.

I should say, parenthetically, in the one case where patient flow analysis wasn't used, the argument was that nonprofits will not exploit market power and, therefore, we should let nonprofits merge. That evidence was subsequently challenged in research publications, and I don't think you can win the day with that argument anymore.

Well, I should add, of course, that based on the structure-conduct-performance studies, the courts have been skeptical about the benefits of competition in healthcare. Combine that with the Elzinga-Hogarty flow analysis and you had failure after failure in the courts.

The flow analysis is kind of silly. It's not tied to any theory. It's worse than structure-conduct-performance, which may be tied to the wrong theory, but at least there's a theory. It's incredibly sensitive to implementation. Depending on where you start, for example, the entire state of California might not be large enough to be a market. You could actually end up expanding the market to accommodate flows and reach the entire state and it's still not big enough. So, like, every hospital in California could merge, and it wouldn't raise price, would be the nonsensical conclusion.

It also conflicts with the empirical evidence, doing retrospective studies. And Cory Capps, who's going to be one of the discussants later, and I did one retrospective study of markets in which there were mergers, markets which would pass muster under Elzinga-Hogarty, and yet those mergers led to statistically significant and substantial price increases. And at the same time, we now have more recent structure-conduct-performance studies suggesting that competition actually does lead to lower prices.

So, how do we get out of this? How do we help the antitrust agencies solve their problems and maybe start winning some cases? Well, here's where structural modeling rises to the rescue, and in the early 2000s, several economists developed structural models to predict merger outcomes. It started with Bob Town, another one of our discussants, and Greg Vistnes, my own work with Cory Capps and Mark Satterthwaite, and a third paper by Marty Gaynor and Bill Vogt. All of us used structural modeling in different ways to reach the same conclusions, which was that, looking at the data, when we applied our models, mergers in markets that would pass muster under Elzinga-Hogarty should not pass muster once you look at these markets more carefully.

What's interesting about all of these papers is

that they're all rather different in their approaches, especially Marty and Bill's approach. The key equations, though, in all three studies are very similar. So, coming at it from different ways, we ended up reaching very similar conclusions because of that. And as I said, the result of our papers is the belief that geographic markets are much smaller than those predicted using Elzinga-Hogarty.

So, the Federal Trade Commission has had all of these things in mind when it set about revising its strategy for assessing hospital mergers. First, it developed some retrospective studies, which is now published, knowing that facts on the ground would trump theoretical predictions. It then challenged one consummated merger, that between Evanston Hospital and Highland Park Hospital, to form the Evanston Northwestern Healthcare System, and they even retained Ken Elzinga, of Elzinga-Hogarty fame, to testify against the use of the methods that he developed for differentiated goods markets, like hospital markets.

And Deborah Haas-Wilson, as the economics expert, developed a theory of two-stage competition in which hospital pricing is determined first by negotiations between hospitals and payers, and then the choice of hospitals made by patients who generally go to

a hospital in their network, an idea that was first developed by Bob Town and Greg Vistnes and then expanded on by myself, Cory Capps, and Mark Satterthwaite.

The FTC won the decision by the administrative law judge, although ENH may have had the last laugh. As a lot of people in this room most assuredly know, for reasons that remain unclear to the rest of us, the FTC chose to allow ENH to remain intact with requirements that the member hospitals bargain independently.

Subsequently, the FTC invited Bob Town to serve as an expert for further merger challenges, and Bob has used a modified CDS model to predict merger effects. The FTC has challenged mergers in Virginia and Ohio. The Virginia merger was abandoned after the FTC challenged it, and I believe the Ohio investigation is still pending -- and, Bob, if I'm wrong, you can correct me on that --

DR. TOWN: We're waiting for a decision.

DR. DRANOVE: That's what I thought. That's pending court decision. I got it right.

And our methods are also finding their way into private litigation, especially in attempted monopolization cases, which are often done in the private sector. A typical example of this is when a hospital that's dominant in a market -- or at least

allegedly so -- goes to an insurer and says, "If you want to have access to us for inpatient care, we want the exclusive for outpatient care," essentially foreclosing the market from freestanding, independent surgical centers.

Hospitals have invariably won these cases when market power was determined using Elzinga-Hogarty. Time and again, seemingly powerful hospitals, to the lay observer, were claiming successfully in the courts that they faced competition from hospitals 50, 100 miles away and, therefore, did not possess market power. But CDS leads to different conclusions and potentially will lead to different outcomes.

So, are we there? Have we finally figured out how to get this right? Well, CDS, we think, yields better predictions than Elzinga-Hogarty. I'm not sure that's been put to the test enough to say that conclusively. It certainly identifies as problematic mergers that have actually led to increased prices, but even CDS can be improved upon.

CDS uses a simple bargaining model. In our model, bargainers are naive. They ignore the impact of their contract on the contracts the insurer will sign with other hospitals. So, if I'm a hospital and I don't have a contract with the insurer, that insurer's

patients are going to have to go to other hospitals, and that insurer is going to have to bargain with other hospitals without me being in the network.

In our model, we ignore that possibility. It was a modeling convenience, but it's just as ad hoc, in its own way, as the Elzinga-Hogarty patient flow analysis, and I think we've done enough other things to make advances, but we still had gotten to a point where we had to make an ad hoc assumption. So, Mark Satterthwaite and I have tried again, along with Andy Sfekas, our third author.

In the immediate aftermath of our paper, we tried to develop a more sophisticated bargaining model, moving into the realm of dynamic bargaining models, but these are brutally hard to develop and solve for symmetric agents, and developing them for asymmetric agents -- not all hospitals are alike -- and doing it in a way that could be taken to the data proved difficult.

How difficult? Well, I don't know if any of you were at the annual health economics conference that took place, I think, four years ago in Madison, Wisconsin, but we presented a version of that paper then, and we've subsequently burned all paper copies and demagnetized any electronic versions.

Our current paper is about bargaining with

foresight, and we relaxed the assumption of naive bargaining, but we only allow bargainers to think one level ahead. So, the insurer and the hospital, say Hospital A, know that if A is excluded from the network, this will affect the outcome of the negotiation between the insurer and Hospital B. And in this way, we're trying to improve the bargaining model, but even this is a hard problem.

There's, I guess, now one published study -- it was just published by Avi Goldfarb, and I can't remember the co-author -- showing how you could identify foresight in real world firms engaged in bargaining. But even that paper doesn't deal with two-sided bargaining of our kind, where you have multiple asymmetric bargainers on one side bargaining with one player on the other. So, we have to simplify it, because it's a devilishly difficult problem.

As I mentioned, to do our model correctly with complete foresight -- and we're focusing on San Diego as a starting point, it's a nice clean market in terms of a reasonable number of hospitals, interesting things -- if you wanted to have full foresight for every bargainer -- if this bargain breaks down, what will happen to the next one, et cetera -- you have to work with 16 million simultaneous equations. Well, we're not doing that.

We're going one step ahead to see if we can document one level of foresight.

We've developed a term that captures that one level of foresight, and just that alone is a contribution to the bargaining literature, because we can now show how you can take data to identify whether people are demonstrating any foresight whatsoever, and we look at several markets in California, and we actually show that bargainers are using foresight; that the naive bargaining model is not sufficient.

We used this to predict the merger outcomes that we had been predicting previously in our older work, and we find that the predicted merger effects are different, because foresight is now being added to the model, and, of course, it's going to be the case. We've got a new predictor of bargaining outcome, so, of course, our predictions are going to change.

But it's not exactly clear to us yet why the predictions are different. Can we tell in which markets the price is going to end up being even higher than our old model predicted and in which they will be even lower? And, Bob, you saw this paper earlier. You should know that we fixed some things. It's no longer systematically overstating or understating. It's now just different, which will be a relief to people who

were kind of hoping that we weren't going to destroy this whole literature. Certainly I would never get retained again.

It's difficult to know, in a complicated model, the role played by simplifying assumptions. So, we make a lot of simplifying assumptions about the bargaining process. And when somebody improves upon our work -- and somebody will, for sure -- the predicted merger effects are likely to change yet again.

So, what should an economist do? Look, economists who are not paid as experts are happy with our results. We've shown that bargainers show foresight, and future research will refine our understanding of how all this works, but economic experts can't wait for future research. You have to testify based on the knowledge at that point in time. So, you have to use the best available models.

Ad hoc approaches and structure-conduct-performance can be touted for their simplicity and durability, but they give poor predictions. We shouldn't use them. Structural models, such as Capps, Dranove, Satterthwaite, are harder to explain, but they give better predictions, and we hope that our new work will give better predictions still.

One advantage to structural models, though, is

that they're easier to criticize. Because of that audit trail, all the potential flaws are there for everyone to see. In addition, every iteration seems to change the results, but that's the nature of the beast.

Over the past decade, with the wars and economic calamity, I've admired the serenity of my cat and the constant joy of my dog, Socrates and Eleanor. To them, the world is a very simple place. If Socrates has his food in his bowl as 8:00 in the morning and at 6:00 at night, he'll sit in your lap the rest of the day, and all Eleanor wants to do is run around outdoors chasing her own tail for hour after hour.

Well, I wish life could be like that for economists, but unlike cats and dogs, we have a conscience. We see more of the complexity of the real world than they do. We realize that some time soon, another veil of complexity is going to be lifted. Things will become clearer. They're not as clear today as they will be in the future. But the view today is clearer than it was yesterday, and antitrust analysis cannot wait for the last researcher to stand on the last pair of shoulders. So, antitrust analysts who work with structural models will never enjoy the serenity and joy of Socrates and Eleanor.

And I have time for questions, I believe. Is

that correct? Great.

DR. O'CONNOR: Are there any questions?

DR. DRANOVE: I've silenced them.

AUDIENCE SPEAKER: You called your dog Socrates?

DR. DRANOVE: Both pets are 11 years old, and by 11 years old, my older son was actually a thoughtful human being. He's 25 now. If anybody goes to Honolulu and just wants to see what bums in Honolulu look like, look up my son. He's actually -- he's got an Internet business that he's making enough money to live in Honolulu without bothering to get real work. What can you do?

That was a tough question, Aviv. I appreciate the -- ah, we have a -- maybe a -- if the question is about my dog, we were supposed to get a boy dog and my son wanted to call him Roosevelt, and so we ended up with Eleanor.

AUDIENCE SPEAKER: Not about your 12 million simultaneous equations, but quality. Of course, focused on price here, and one issue that is frequently raised by merging parties that they claim is -- should be taken into account is improvements in quality, and measuring the welfare effects of quality changes can be very difficult.

Have you given any thought to quality issues?

DR. DRANOVE: I've actually not personally written any papers on mergers and quality, but there have been some studies published, and Cory or Bob -- Bob, you probably know this, because you must have written part of this in your review. The evidence seems to be rather mixed as to whether mergers lead to improvements in quality overall. So, I think the argument in the courts is always one side making strong claims and the other side saying -- you know, everybody's going to make strong claims, but the overall evidence is mixed.

Well, in that case, since it's 1:40, I guess we have managed to get ourselves right back on schedule. Let me introduce the first speaker from the University of Toronto, Matt Grennan, who will talk about discrimination in bargaining, empirical evidence from medical devices. And his discussant will be Bob Town, from The Wharton School at the University of Pennsylvania.

PAPER SESSION TWO:

ECONOMICS OF ANTITRUST

DR. GRENNAN: Okay. So, this is a paper about the prices being put on a certain medical technology, called the coronary stent, and I apologize that this is after lunch, but we had a relatively healthy lunch here.

So, the coronary stent is a little mesh metal tube that is implanted in to keep arteries propped open after they go in with a catheter to try and clear blockages in the arteries around your heart.

This is sold in -- from device manufacturers who make these stents to hospitals to implant them in a procedure called an angioplasty, and the hospitals generate revenues by performing these procedures, and they're reimbursed by Medicare or private insurers, but the cost of these stents comes directly out of the revenues that they're generating.

And an interesting thing about this market is, like many business-to-business markets, prices are negotiated. So, we're kind of building on a theme here, I guess. And, interestingly, different hospitals pay very different prices for the exact same stent from the same manufacturer.

So, this picture may look similar to one you saw in the previous presentation. So, this is a histogram

of the prices being paid for a particular coronary stent by different hospitals in the exact same month.

So, as you can see, on average, this stent is selling for about $2,400 each, but some hospitals are paying $1,800 apiece; some hospitals are paying more than $3,000 apiece. And these numbers add up, because the average-size hospital implants about a thousand of these stents per year. So, the difference between being in the 25th or the 75th percentile in this price distribution is about $300,000 per year, say four nurses' salaries.

Now, the question I'm going to look at is, what happens when you look at cases where prices have to be more uniform? And this sort of question is raised a lot in healthcare, because -- well, one, look at mergers. So, mergers are a case where two hospitals used to be two different entities; they received two prices. They merge, they're now one, and you're setting the same price to this merged entity, except for in the case of Northwestern and Evanston, I guess.

And also, group purchasing organizations. So, group purchasing organizations are third parties that -- in the healthcare world, they play a big role in purchasing for a lot of goods, especially the sort of things like, you know, bandages, syringes, these sorts

of products. And, interestingly, they don't play a big role for things that get called physician preference items, things like coronary stents, like defibrillators, Pacemakers, these sort of products. And hopefully some of my analysis here will be able to maybe suggest a few of the reasons why that may be.

There's also been a bit of discussion recently on transparency of prices in healthcare in general, and to the extent that we think that transparency is the one of the things that enables setting very different prices or ending up negotiating very different prices with different buyers, we might think that will affect the amount of uniformity in the prices in this market.

And now, so, there's been talk and evidence on, you know, the effects of some of these things, but it's just really not clear in terms of, you know, mergers, for example. It seems the answer is usually it depends. Does a merger generate cost synergies or lower costs for a hospital, in particular for things like inputs, like coronary stents? I told you about the group purchasing kind of conundrum in terms of how it seems to work for some types of products and not for others. And I think the price transparency thing is just kind of a general open question.

Now, this is a market that's a little bit

different than the one we saw earlier. So, this is not a market where we're thinking about search and these sort of things. So, there are four -- at least in the U.S., there are four device manufacturers selling these coronary stents, and there are -- you think a given hospital is contracting usually with all of these four device manufacturers. And there are 11 differentiated types of stents that these four -- so, 11 in total, so each manufacturer offers about two or three different stents that are a differentiated product.

And they're differentiated in the sense of some are going to be better for particular patients, and they're also differentiated in the sort of usability factor where, you know, some doctors may find some kind of more appropriate for their tastes than others.

And so the set of things that we want to think about, we almost need to move back a step and think about, you know, research question zero, if you will, is where is this price differential coming from in the first place, right? So, one story could be it's all demand. So, different hospitals, you have different groups of physicians who have different tastes over these products or treating patients with -- a group of patients with different characteristics, and all this price dispersion would be due to this heterogeneity in

demand across markets.

In that world, we're squarely in the world of price discrimination with oligopoly, and what the theory there would say is that what's going to matter is, what's the nature of that heterogeneity in demand? So, is it that there is kind of vertical differentiation in the sense of you have some hospitals who generate more revenues for each procedure and so have a sense of higher willingness to pay for all of these stents, or is there a more horizontal differentiation factor where you have, you know, one hospital in Evanston where they're facing a population of patients where, you know, stent A is usually the right stent to use or often the right stent to use in some, you know, other hospital in Joplin, Missouri, where, you know, stent B is usually the right stent to use?

So, it's going to turn out that in the first case, so in the vertical differentiation case, what we would actually see is usually competition would intensify if we went to more uniform pricing -- so, it would be good, in general, for lowering prices -- whereas in the horizontal differentiation world, competition will actually tend to soften when we move to more uniform prices, because it allows each manufacturer -- they will just set -- they are going to

be setting one price, and that price is going to be set to extract surplus from the hospital, in this case, where their stent is more valued. So, that's one potential explanation.

Another one is it's all negotiation. So, there's no differences in demand across these hospitals, but we know in markets where prices are negotiated, that's only going to delineate a range in which prices could fall, and it's all about bargaining within that range, and you just have some hospitals where they are just good at negotiating a better deal, conditional upon demand.

And so the goal of this paper is going to be to kind of tease these two things apart and then look at changes to more uniform pricing, such as a group purchasing organization or hospital merger, and think about. How are those different sources of heterogeneity going to play into the prices we would see under those counterfactual scenarios?

So, there's a small number of papers that I cite here at the bottom, and that's close to the world of papers that are dealing with this, especially on the empirical side. As you know, there is huge literature on both price discrimination and bargaining in the theory world, but on the empirical side, there is much

fewer, and I don't think it's been driven by the fact that these aren't important economic phenomena that we want to be looking into. It's usually driven by the fact that it's been very hard to find data on a lot of these situations.

So, this study is going to be based on a data set that actually has data on the prices and quantities being exchanged between the stent manufacturers and different hospitals. What it is, is it's a survey of 96 different U.S. hospitals monthly, so it will say for the -- so, it's -- for each stent, observe how many that hospital uses and the price that hospital is currently paying on a monthly basis for 3 1/2 years. So, it's over 10,000 -- the unit of observation is a stent in a hospital in a given month, and the data here is just to give you an idea of what it looks like and also give you a bit of a feel for the market.

So, in observation, I'll see the year; the month; the hospital is an anonymous identifier in this data -- which is one of the kind of unfortunate things about this data, where it kind of lacks detail; the product -- so there are nine different what are called bare metal stents and two different drug-eluting stents on the market; the manufacturer who makes that product; the quantity currently being purchased; and the price

that it's currently selling for.

I told you I don't see a lot about the hospitals. Fortunately, what I do see is -- the number at the end of that table there is I do see the number of diagnostic procedures, diagnostic angiographies that hospital does each month, which gives me a nice idea of the set of patients who could possibly be getting a coronary stent in that hospital.

Now, you noticed in that data set I showed you, there was no number for willingness to pay or the cost to manufacturers or some sort of measure of bargaining ability, as I'll call it, and the way I'm going to try and get at those numbers is by developing a model that's going to, I hope, capture the features of this market and allow me to infer those parameters.

So, this model is going to have two stages. So, the first is a pricing model. So, prices are negotiated but also incorporating competition between these differentiated products. And in the second part of the model, once these negotiated contracts are in place, patients will arrive and physicians will make decisions about which stent to use to treat that patient.

So, the demand side I'm not going to say too much about, because in many ways, it's, you know, very standard, building on, you know, a large, large

literature on demand estimation. The important things to point out here is it's really important to allow for heterogeneity in demand across hospitals, which, thankfully, due to the panel data, I can do very flexibly, but also heterogeneity across patients and doctors within hospitals, so being very flexible in the shape of the demand curves at the hospital level.

I'll do that with a random coefficients discrete choice model, and what I do think is worth talking about just for a minute, though, is that there are actually two kind of interesting new sources of identification for demand in this case, where prices are negotiated and fixed over a period of time in longer term contracts.

So, to give you some kind of intuition, this is kind of the Econ 101 version of a way in which this identification argument is different or can be different. So. Just think of demand and supply curves. There's a pricing equilibrium, and demand shifts over time. So, this would be -- the idea is, you know, physicians -- there's turnover at the hospital level in the physicians who are there, there are new studies that come out that suggest one stent maybe is better than we thought it was before for a certain population of patients, and as you'll see, because prices are fixed in a long-term contract, price is kind of out of

equilibrium in a sense now.

Well, assume for a moment that demand doesn't shift in this next period, but price comes up for renegotiation. Well, the move is going to be exactly along the demand curve, helping to identify demand. The way I'm going to operationalize this is I'm going to use lagged prices for the same stent, stent j, at the same hospital, h, in month T minus 1. So, all of the identification here is going to rely on sort of within a stent hospital, looking at variation over time to identify everything, and then I'm going to lag that to allow for the fact that, you know, when prices were renegotiated, there could have been a shift in demand.

And I'm also going to use -- so, this other new source of identification is a bargaining ability. So, in a model where prices are negotiated, the bargaining weights that determine that split actually give us another supply shifter, and I am going to try and capture that by using the average prices of other stents at that same hospital in the previous month.

Because these negotiations don't all reoccur necessarily at once, you could see new contracts being negotiated by other stents. If the bargaining of that hospital has shifted, these average prices should be capturing that.

All right. So, on to the pricing side of things. So, the idea here is to have a model that's going to incorporate, you know, costs, demand, and competition, how those feed into prices, but also allowing for bargaining in that there's a range of potential prices that could be agreed upon. And I am going to do this using a model that's a Nash equilibrium of bilateral Nash bargaining problems.

This builds on a kind of -- it's related to an area of theory of negotiations in -- with externalities, and as -- similar to the models that are being used in some of -- some other empirical applications where prices are negotiated.

And this may -- you know, so it's -- it may look a little different at first until you realize that it's really just a generalization of the model we know and love, where manufacturers set prices to maximize manufacturer profits in a Nash equilibrium. It's just instead of maximizing manufacturer profits, it's maximizing the bilateral Nash product, where it's manufacturer profits raised to some parameter which is going to capture their bargaining ability, or how -- the weight the manufacturer profits have in the function that's being maximized, times hospital surplus, minus the disagreement point of the hospital, of what happens

in the case when a negotiation with that particular stent would break down, and that's the point where some of these different assumptions that David was talking about actually come into play and what that disagreement point ought to look like.

I am going to assume it's a sort of Nash assumption, a more static assumption, where all the other prices would remain the same in the case where a negotiation were to break down. And then that would be raised to a parameter that would give you the hospital's weight. So, importantly, the reason this is a generalization of the model we know and love, if the hospital's bargaining parameter there is zero, then we're in the world of setting prices in a Bertrand-Nash equilibrium. If the weight on the manufacturer profits is zero, then its prices are set at marginal cost.

So, I think another way to think about this model, which I think makes it clearer, in a way, is to look at the pricing equation that pops out of those first-order conditions. So, you have price as cost plus a margin, and that margin is the ratio of the manufacturer's bargaining weight to the whole of the manufacturer and hospital. So, this is going to be a fraction between zero and one, and that's going to be times the entire surplus that's up for negotiation.

That surplus has a part -- if you look over on the right-hand side there, this is just the -- you can think of it as the preunit added value of a product, right? So, the extra surplus is being created using that product compared to if that product were unavailable. But the part of that surplus that's going to the hospital actually needs to be adjusted by this elasticity term here, where because quantities could -- and I find a little bit do -- depend on prices, then increasing price by a dollar doesn't mean another dollar in the manufacturer's pocket, because we're moving along a curve and quantities are going to decrease slightly.

So, the way I'm going to operationalize this in an empirical setting is to say that the price of a given stent, j, in hospital h, at time t, is going to be equal to a cost term, as a parameter to be estimated, plus a term that's this added value term, which I can calculate from the demand side estimates, times this ratio of bargaining abilities.

And the decision that needs to be made here -- so, we can identify -- you can see clearly there how you identify bargaining ability separate from costs and that as added value changes, to the extent that price changes, that's going to help us figure out what bargaining abilities are, the correlation between those

two terms, whereas cost is more of an intercept term in this equation. But the one thing that we do need to take a stand upon is that, you know, trying to estimate a full distribution of different costs and a full distribution of these bargaining abilities would be estimating a separate intercept and slope parameter for every observation in a regression.

So, in this case, there's not a lot of reason to think that costs are going to be varying a lot across hospitals or across time, so all of the unobservable portion is going to be going into bargaining abilities in this case. So, it's just something to keep in mind as you look at the results.

So, this is a table just summarizing all of the results from the demand and supply side estimates. So, if you look down the rows, so this is listing each stent available on the market in this particular month of September 2005, and on the left-hand column is the price data. So, it shows you as both a mean price and a standard deviation of price across different hospitals in that month.

The next column is the cost estimates. So, one thing that ends up being -- so, the costs in this exercise, in the paper I compare them to the costs you'd get if you kind of, you know, didn't use the bargaining

model. They're much better than that, but unfortunately, they're actually not very tightly identified, and the reason why is actually if you look over at the added value estimates, which are the last column -- so, this is kind of the extra surplus that's being created, on average, for the patients that are treated with that stent versus if that stent weren't available -- and they're big.

Even when you look at the standard deviations across hospitals, even at the hospitals that have a very low added value, they're still a pretty big estimate, and if you remember -- so, you know, price is -- cost is an intercept term, plus the bargaining, times the added value. What this means is we don't have a lot of data down near the intercept, right? So, cost is going to be tightly pinned down by observations where added value is close to zero, and we're just not seeing that in the data here, which means that the costs aren't going to be very precisely estimated here.

The fortunate, I guess, side of this is that the subsequent estimates are not qualitatively at all changed if you were to move around costs, you know, into every reasonable end of the spectrum of costs that you might think might apply, and even quantitatively, the changes are rather small. And that's an appendix in the

paper.

So, the other interesting column that pops out here is these bargaining ratio estimates, which are the column that I haven't spoken to yet, and you'll see there that you also have -- you know, on average, it looks like manufacturers are getting about a third of the surplus that's being created, but there's quite a bit of heterogeneity in those standard deviations across hospitals.

So, back to, where is that heterogeneity in prices coming from? Part of it's coming from this heterogeneity in these added values, which is demand side heterogeneity, and part of it is coming from heterogeneity in these bargaining abilities.

Now, one thing that, you know, kind of pops out of this, so, this bargaining ability parameter seems to matter a fair amount in this market. Where does it come from? So, you know, back to the -- there's not a lot of specifics we know about the hospitals. There's not a lot I can tell you about that, other than the fact that, you know, having this panel data actually helps me. At least I can look at, you know, what seems to be firm-specific about it.

So, you can regress that term on firm dummy variables, has R2 of about 0.41, but more importantly,

for this study, it actually -- the coefficients on those dummy variables give me an idea of a firm bargaining ability for each firm, which I am going to use in the counterfactuals. So, for each manufacturer, I have a bargaining ability for them, for each hospital a bargaining ability, and that distribution of different bargaining abilities across hospitals is going to be helpful in thinking about, you know, what might be the bargaining ability of a group purchasing organization or a merged group of hospitals?

So, in the counterfactuals, what I look at is cases where the same sort of bargaining specification, but now demand is merged across a group of hospitals. So, in the GPO case, looking at all hospitals in my data set as if they were one big group purchasing organization. In the case of some mergers that I simulate, looking at randomly selected groups of seven hospitals.

And -- or, again, remember this -- so, this beta or this BH parameter is going to allow for the hospitals to bargain collectively at this bargaining weight, and remember, again, the case where this is zero is the one where manufacturers set prices.

So, this graph captures the case when we treat all the hospitals in the data set as being part of one

big group purchasing organization, and where it maps out is the change in hospital surplus on the Y axis for the bargaining ability assumed of that hospital group on the X axis. So, the bargaining ability is done as a ratio to the average hospital bargaining ability.

So, you'll see in the middle there a one. It's kind of like fixing the change in bargaining ability and really isolating what's the change due to this price discrimination effect of the change in competition when we move to this group purchasing organization.

And this fact that prices tend to increase, which decreases hospital surplus, is exactly pointing to this fact that in this market, what we actually have is more this horizontal differentiation across hospitals, where different hospitals prefer different stents, which is leading to price actually becoming -- price competition becoming softer, or less intense, when we move to a more uniform pricing world.

Now, interestingly, you can think about -- so, what -- you know, what would be the bargaining weight that's needed to overcome that? And it turns out if you map up to where you cross back over to hospital surplus becoming positive there, so you'd have to have a bargaining weight over the 70th percentile of all the hospitals I observe in the data set in order to overcome

this.

So, I think this is where I was telling you, you know, offers one potential explanation for why would you see group purchasing organizations playing a big role for more products where you don't have this big horizontal differentiation in demand and perhaps not for the case of these physician preference items, like coronary stents.

So, I'm running out of time here. So, I'm actually going to skip the merger section, but if you're interested, there's plenty of it in the paper. The most interesting thing about it is that it actually -- you know, these different hospitals, different groups of hospitals that I looked at mergers between, will have different levels of symmetry or asymmetry in their demand. So, different amounts of horizontal or vertical differentiation in their demand, which really helps kind of put a number on this competitive effect and how it might vary as the group of merging hospitals varies in their demand asymmetry.

So, you know, I think the take-aways here is, you know, put some numbers to a lot of things that have been discussed in theory and in a context where they're important and of interest, and I think, you know, for further research, I mean, I think one of the things that

it leaves me wanting to know more about is what's behind this bargaining ability, you know, and just anecdotally, in talking to people in the market, you know, these are parameters, these sort of Nash parameters which we usually think of, you know, coming from somewhere -- we're not sure, as economists -- but it actually seems like there's some economics underlying these, right?

So, there are things like human capital, the people involved in negotiations, the way organizational incentive structures are set up, that are determining, you know, what are these bargaining weights that end up being the firm's, you know, Nash parameter? And this other one, you know, it's obvious that this is a very static analysis in some sense, you know, especially in a market like medical devices. What I think we would like to move towards is something that's a bit more dynamic and looks at, you know, entry of these products and how it affects innovation and access to technologies.

Thank you.

(Applause.)

DR. TOWN: So, I've known this paper for a long time and admired it for quite a long time, and Matt's been working on it for a while, and it's very polished, and it's really a great paper. It really is one of the better IO papers in healthcare in the last five years,

so I highly recommend it to you.

Imitation is the sincerest form of flattery, and I have a paper with Aviv and Chris Garmon, who I saw here, where we're basically applying some of the ideas in this paper to hospital bargaining, as Dave talked about. So, it's really a great paper.

The reason it's a great paper, it's an important question, like why is there price variation? What's going on there? The data is excellent. The methods are appropriate. He's developed a bargaining model framework and applies it to the data in a very sensible way. The counterfactuals are interesting that he performs in the paper.

He didn't really get a chance to go into it in much detail here, but they're really quite interesting. He did one about what happens when you get uniform price versus you allow for negotiated prices, and then, also, simulates hospital mergers and what the impact of, in some sense, increased monopsony power has on prices. And there's not much work on monopsony power and simulating monopsony power, and that's a pretty interesting idea.

And the market he's looking at is stents, and he didn't really talk about that that much, but it's really an important market. It's probably the most important

medical device market there is. It's a 5 billion, 6 billion dollar industry. It's huge. There's a lot of innovation going on. I'll talk about that in a sec.

And then more importantly, there is not a lot of work done in the medical device industry, say, relative to pharma, but it's an industry that's approaching pharma in size and importance, and pharma seems to be in decline. So, it's an area that I think deserves a lot more attention.

So, I want to talk about -- now that I've talked about what I like about the paper, now I want to talk about what I think -- areas that need to be improved upon in future work, and one is this Nash bargaining issue. And David alluded to it in his talk, but here, the assumption is that when I negotiate -- I'm a stent manufacturer, I'm negotiating with a hospital, I don't know what's going on in any other negotiations that are going on, and I assume that my outcome is just this unilateral bargaining, and any other outcome is irrelevant to me, but that precludes exclusive dealing in an equilibrium, I think.

So, you can't do any deal where you will be my exclusive manufacturer, which could be important here. And it will feed back, in effect, inferences about bargaining ability later on, so the model may be a

little bit -- may be misspecified, if that's important.

Also, in this market, it's not that just hospitals and manufacturers negotiate prices individually. There actually is a list price, which is kind of -- there is option value to go off the list price, and at least anecdotally, drug-eluting stents, they were mostly on list price for a while before there was more entry, but...

God, this is a messy slide. So, the issue is, what's the right measure of surplus? So, in the model, he has kind of this patient-physician utility and welfare from the exchange, and I'm not sure that's the right way to think about it. I think it's really kind of what the revenue is to the hospital from the deal and some value that the patient gets, kind of merged together, and in the framework Matt put up there, it's really just, you know, the utility. And so there's some -- there might be some disconnect there.

It could be important when you think about reimbursement, because there's no -- because reimbursement kind of doesn't enter into his model, like if you change, you know, CMBS reimbursement rates, they really should have at least no direct effect in this framework, but it seems likely it would have an effect on both the bargaining leverage and the money at stake.

So, that's where it might be important.

And also, that these products are bundled. They're not just negotiated over stents. Boston Scientific, which is -- makes the Taxus stent, and J&J, which has the Cypher stent, or at least used to, they sell a whole range of medical products to hospitals. So, they're negotiating over a whole range of products.

And in the model, bargaining skill is very -- skill varies quite a bit, and that's because it's a residual in the framework, and residuals are often -- you know, the data is hard to explain, the prices. And so that strikes me as unattractive. And furthermore -- and Matt alluded to it -- it's difficult to perform counterfactuals without making some ad hoc assumptions about what happens to bargaining skill under different scenarios.

So, here's some ideas, what I think -- where you can take this. One is -- and Matt alluded to it -- is entry and exit. It's really important here -- here is the -- so, Matt's kind of data kind of goes up to here, roughly, but then there is entry from Medtronic, that's the Endeavor; Abbott, which is -- actually, it wasn't entry, it was purchase; and then here's the Taxus. So, there's been a lot of interesting dynamics, and there's -- these stents are -- two of these are owned by

the same company, I think Chromas and Taxus are -- anyway, I get confused, but two of them are owned by the same guys. And then there's Exit.

So, J&J, which is the initial -- they were the first folks to bring the stent to market, they -- the stent story is really interesting, but -- you know, it was basically developed by this doc in a garage, and he eventually sold it to J&J, and they brought it to market in '94 with the bare metal stent. And then there was innovation where the drug-eluting stent came to market, but ultimately, J&J, which was the first in the market, ended up exiting, and it's kind of interesting how they became the first and were the innovator and ultimately had to exit. And part of the reason they exited were they screwed up.

And then there's mergers in the device industry that are important, that you could use this as an input into analyzing. In particular, there is a pretty well-known Boston Scientific/Guidant-Abbott merger, in that Boston Scientific bought Guidant, but they had to divest their stent business to Abbott, and so there you could simulate sort of what would be the right outcome here under different scenarios.

And then J&J bought Conor, which also was a stent manufacturer, but I guess their stent blew up and

wasn't good, and J&J lost a lot of money, but there was a lot of interesting overlap issues there. So, this is, you know, a framework where you can address these kind of complicated merger analyses in an interesting and realistic way. And this is kind of what happened in the Boston Scientific and Guidant deal, and Boston Scientific kind of ended up taking it in the shorts.

And then finally -- and I talk about this idea mostly because I'm working on it with Matt, so I can do a lot of preadvertising before we get anywhere on it -- is that there's a lot of variation across the globe in medical device regulation, and that's pretty interesting, to see how that impacts consumers. And we can look at the variation between the EU, which is actually a very lax regulatory environment for devices. It's very easy to bring a device to market in Europe, where in the U.S., it's difficult -- and perhaps more difficult recently -- and firms complain a lot about the FDA and their inability to get products through the FDA, versus Japan, which is really hard to get into.

And so we can look at how entry of devices across these different geographic areas impacts perhaps health and consumer surplus and perhaps safety of the devices. And, generally, when you talk to device manufacturers, they say, "Well, we just go to Europe

first, because we can get into Europe, we get data, and maybe we'll come to the U.S." So, you see a lot more devices and a lot greater breadth of devices in Europe than you do in the U.S., and so this could be used as an input into a model that starts to quantify these impacts.

I think that's where I'm going to stop. It's really a terrific paper, and I highly recommend it to you.

(Applause.)

DR. DRANOVE: The official timekeeper who's got those signs that say it's time to quit seems to have disappeared. I know we're running a little behind. Matt, you can respond to Bob privately. One or two questions, possibly, for Matt, but no more than that. Anyone?

(No response.)

DR. DRANOVE: Okay. Then our next paper will be presented by Sonia Jaffe from Harvard University. She's going to present the first-order approach to merger analysis, and Cory Capps, from Bates White, will be the discussant.

MS. JAFFE: Thank you.

Thanks for having me. This is work that I am working on with Glen Weyl, who -- I don't know if he's

here yet, but he will be around tomorrow. And, basically, we're trying to build on a growing literature in the merger analysis arena of taking what we call the first-order approach to merger analysis.

And so there is some work by Greg Werden that shows that you can calculate the hypothetical efficiency gains necessary to offset the pressure to increase prices and then work by Joe and Carl developing the upward pricing pressure, which then got implemented both in the U.K. and in the most recent U.S. Merger Guidelines. And we're very much building on this literature.

The first thing we do is we generalize upward pricing pressure. We want to allow for nonpricing conduct, so not assuming that it's always a Bertrand-Nash equilibrium, and this generates two changes to the formula. One is a more general diversion ratio, and the other is what we call an end of accommodating reactions term, and I'll explain both of those.

And then the second thing we do is that we formulate what we call the merger pass-through rate. Because upward pricing pressure is really just pressure, it's a change in incentives, it's not measured in units of price, in order to get an estimate for how much

prices are going to change, you need to know to what extent these pressures, these costs, get passed through by the firms to prices. And it's a combination of what would be the premerger and the postmerger pass-through rates.

And then the last thing we do is we weight these by quantities in order to get an estimate in the change in consumer surplus.

So, starting off with just a real general model, you've got n firms, they can be multiproduct firms, and they have a strategy dimension, which is the same dimension as the number of products that they sell, and they've got your pretty standard profit function, which depends on the prices that are generated in equilibrium when all the different firms play their strategies.

And because we don't want to just limit ourselves to Nash equilibrium, we're going to think about total derivatives, which means when a firm changes their strategy, there's the direct effect on prices and thereby on quantities, but there's also indirect effects, which is when a firm changes its strategy, it may be that other firms change their strategies in response. And combining those two, you get the total effect of a firm changing its strategy.

And in order to kind of keep this in the

existing framework and because we're mainly focusing on -- we're focusing on how these things are going to affect prices, it's a lot easier to think of these strategies as prices. And this isn't to say that firms are necessarily setting prices, because then you would just be back in a price setting, but it's just saying that for most strategies that you can think of firms playing, there's going to be an invertable map from those strategies to prices. And so, therefore, you can think of the merging firms as choosing prices, but this is going to change the reaction functions of other firms.

And the clearest example to see this is if you think firms are actually playing a Cournot equilibrium, but you want to talk about their behavior in terms of them setting prices, then when one firm raises its price, all the other firms are going to react by raising their prices in order to keep their quantities fixed, because that's what Cournot does. It keeps the other firms' quantities fixed when one firm changes its behavior. And so we can incorporate those reactions into these total derivatives that we're thinking about.

And so you have your very standard premerger first-order condition, and then once the firms merge, these things change, because firm j is no longer

reacting. So, you can see that now, when you look at the total derivative, the effect of the change in i's price has on all but i and j's prices, in order to get the total derivative. And this is going to lead to a different diversion matrix, because you have to think about what the diversion ratio would be if firm i raised its price and all the other firms were allowed to react but firm j was not allowed to react. In doing this, you can get the new first-order condition, and you subtract the two, and you get the generalized pricing pressure.

And so as I mentioned, this diversion ratio is going to be different from the one in a Bertrand-Nash in prices model. And then there's also this end of accommodating reactions term, and that comes from the fact that before, when firm i raised its price or changed its strategy, the other firms would react, which could frequently have a dampening effect on the change in demand. And so postmerger, in general, you're going to expect that the merging firm's demand is in some sense more elastic, because the firm j is no longer reacting.

And in general, these things are going to go in the opposite direction. The diversion ratio is going to be greater, but you're going to be subtracting off a positive term.

So, one thing that we look at a little bit in the paper and that we're interested in exploring more is how different we really expect this to be from Bertrand competition, but there are some circumstances where this end of the accommodating reactions term is really going to make a big difference.

And it's important, once you combine this with the pass-through, that you use the rates from the same model. If you're using a Bertrand model to calculate the pass-through rates, then you really need to use the Bertrand UPP, but if you're thinking about a Cournot model when you're calculating or estimating the pass-through rates, then you need to have the corresponding generalized pricing pressure.

Just to explain this briefly, as I said, in Bertrand, this is exactly the same as UPP. With multiproducts, you just use the matrix inversion. For quantities, as I mentioned, this is essentially the other firms reacting to a change in i's price so as to keep their quantities fixed, and then you, in turn, see what effect that has on firm i's quantity.

So, this is the more general pricing pressure, and then we want to convert this into actual changes in pricing, which is where we get to pass-through. And this has come up in the literature before, thinking that

because these are essentially opportunity costs of sales, that something like pass-through should be used to convert them into price changes, where pass-through specifically refers to when you have marginal physical costs, how those convert into price changes.

But there's been some disagreement as to whether you want to look at the postmerger or the premerger pass-through rates. And just to try to give a bit of an intuition of what these pass-through rates are and why they're what matters, this is my drawn-in-paint graph here, but if you look at the premerger profit function with price on the X axis, premerger, you're going to go to the maximization point. Postmerger, you have now combined two. So, the scales on these might be different, but whatever the slope of your profit function is at the premerger price, that's the pricing pressure. That's the UPP or the GePP.

But then to know how that converts into a change in price, you have to know what the curvature of this profit function is, because if you have the red line, which is highly curved, you're going to get a very small change in price, because that change in slope very quickly tailors off. If you have the blue curve, which is much less curved, then you're going to get a much larger change in price.

And because the pass-through is related to the inverse of the second derivative of the profit function, the higher the pass-through, the less curved the profit function is. And so that's why we're looking at pass-through rates, to try and figure out how much prices are going to change for a given shift, a given angular shift in the first-order condition.

So, we use a Taylor expansion, so this is a first-order approximation to what the pass-through rate is going to be, and we show that it's this thing here, where f is the premerger first-order condition and g is the pricing pressure, so that f plus g is the postmerger first-order condition. And you take the derivative of that expected price and take the inverse, and that's your pass-through rate.

And so this is related to both the pre- and postmerger pass-through rates. So, here, premerger, as I said, it's just the inverse of the second derivative of profits, which is the first derivative of the first-order condition. Postmerger, you multiply that by this matrix which has to do with diversion ratios, and the reason is that postmerger, the marginal costs of one firm enter the first-order condition of -- the products that did belong to one firm enter the first-order conditions of the products that did belong to the other

firm.

And so this is basically confusing the postmerger curvature, as I showed in that picture, but not using this diversion matrix, which kind of distributes the costs to the different first-order conditions. And the reason is that these aren't physical costs. These aren't marginal costs of selling a different product which are going to enter in different places depending on how demand for a product is affected by another product's price. They are pure opportunity costs, and so they enter linearly in each of the first-order conditions, which is why you just have the curvature matrix, not multiplied by this matrix based on the diversion ratios.

And so then the question is -- so, that's kind of the theory of why -- how to get at the main things that are important to how much prices are going to change. And then the question is, what do you do with it? And we can't observe merger pass-through before the merger happens. If you have some very restrictive scenarios, you can sometimes calculate it based off of the premerger first and second derivatives of demand, which you could get if you had premerger demand and pass-through rates, but generally, you're going to need some other assumptions.

And we've done a little bit trying to see -- basically, the idea is that if this pricing pressure is reasonably small, then you would expect that its derivative is also small and that the merger pass-through rate will be well approximated by the premerger pass-through rate. And so using the premerger pass-through rate is going to be a reasonable -- in that case will be a reasonable approximation for the merger pass-through rate.

Another concern is the error terms, since this is an approximation. Here, you know that the first-order conditions are not highly curved or -- and we think that smoothness is generally a reasonable approximation for most demand systems and that if the pressure -- and then the other issue is the smallness, and if the pressure is very large, then it seems like the merger is a bad idea and you're less worried about getting the estimate exactly. Yes. So, that's the two parts of predicting price changes.

And then to get welfare changes, we don't have a strong opinion on the whole consumer surplus versus social surplus debate, but either way, if you have the elements to calculate these price changes, then you can calculate both the change in consumer surplus and the change in social -- the change in dead weight loss or

the change in social surplus and use that to create some kind of index of the percent change in surplus in the market.

And this could also be useful if there is nonprice concerns -- first of all, if there's fixed costs, those would be added in here, and then if there is other nonpricing concerns in the market, putting them all into dollar terms is what we think would be the easiest way to compare across the different effects of the merger.

So, in one sense, this is a very approximate approach. It's not giving a whole structural model with very specific parameters, but in some sense, we see that as a benefit, because it's more flexible in that sense, right? If you took all of the assumptions that went into a specific demand form, if we took that structure, you would get the exact same answer. You could plug it in here, add some higher order terms, you would get the exact same answer.

But we think that sometimes by assuming a functional form, you're making a whole bunch of assumptions about pass-through rate and about other factors that really matter without realizing -- sometimes without realizing it and certainly without acknowledging it, and by putting the things that really

have the direct effect up front, we think that it's sometimes easier to really see what assumptions you're making and what assumptions are driving the conclusions.

And we're not saying that this is necessarily -- like, should completely replace structural estimation. You may want to do structural estimation to get some of the demand derivatives, diversion ratios that we use, but we think that rather than tying down the pass-through rate in a functional form assumption kind of at the back end, it's better to have it up front, where it's easy to see what's going on.

And we also think -- you are going -- frequently, to a point, you'll still need to make some simplifying assumptions, either about conduct or about the level of heterogeneity or about the pass-through rates, but I think that if you have this framework, then you can very easily test those different assumptions, see which are driving them, see whether the assumptions are reasonable, and not have them be kind of locked down in functional forms or in things that are not at the forefront.

And just as we emphasized, we think it's really -- just to do robustness checks, it's much easier really to just try a lot of different things. If you have uncertainty, rather -- because a lot of Glen's work

has shown that traditional functional form assumptions really tie down the pass-through rate in ways that you often don't realize.

So, that's pretty much it. We generalize the pricing pressure and show how to convert it into prices using pass-through rates, and, in general, this approach of using pass-through is really saying that even though mergers are a distinct change in some sense, you can use these methods of sufficient statistics and small changes to approximate their effects.

And we think there's a lot more work to be done in this area, figuring out how good an approximation this is, when it works well, when it doesn't, possibly add dynamics or other richness, possibly about quality, and, again, thinking about the best ways to simplify the formula for applications.

Thank you.

(Applause.)

DR. CAPPS: All right, thanks.

This was a good paper or a good addition to a series of papers that's coming out that I think does two really nice things. It helps sharpen our thinking about what could be done in merger analysis in the real world and what should be done, and if we're really lucky, there will be at least some intersection between the

two. So, even if this isn't a tool literally used for screening mergers or at trial -- and maybe there's a role for that -- it is a tool at least for improving the way we think about what we do and maybe the limits and lower bounds on the usefulness of things like UPP.

So, this builds on some prior work and says, you know, these prior analyses, going back to Werden's 1996 paper, but maybe more popularized by Farrell and Shapiro in various papers on UPP, it says, you know, we can really get an idea of the gist of what a merger effect will be by thinking about the impact on the pricing incentives of the merged firms when the firms are producing substitutes.

And the intuition I think for that is pretty compelling and clear and actually easy to explain, which is that a merger of substitutes will reduce the cost of increasing price, or as they word it in the paper backwards, sort of will lower the -- raise the benefit of -- no, lower the benefit of cutting price, but it's fairly symmetric.

And it's important to point out, you know, it is a lot of work in terms of the equations. There really is a benefit and an improvement over what is reflected in UPP, because it captures real world economic phenomena that are omitted from UPP. The real world

phenomena are not the misplaced dots that are supposed to be at intersections of things.

But the idea here is that when you have strategic complements, if a firm increases its price, that's going to induce other firms to also increase their price, and that's kind of a freebie for the first firm that's increasing its price, right, because that means it doesn't lose as much quantity as it otherwise would. And that type of phenomena really is nowhere captured in the basic UPP formula of diversion times margin of the merging partner's product.

But in particular, I think it's also a nice insight, because it runs kind of counter to your intuition, which is that you think -- if you think just about the unilateral incentives of the merging parties, you might underestimate the merger effects, because if there's strategic complementarity -- do I need to stand here? Thank you. I saw someone moving their head and I thought I was saying horribly wrong things, which may be true, but now I know it's just the microphone.

Normally in cases of strategic complements, you might think that if you just think, what is the incentive on the merging parties to increase price, and you say that's positive, well, the real world effect would be even larger if that causes a price increase

that is then met by an accommodating response and you shift the equilibrium. So, I do think this is a nice advancement that captures some intuition that's left out of UPP.

But then now I'm in economic consulting and litigation, and this leads me to the next step of thinking -- and I also used to work at DOJ -- how would this actually play out in real world merger review, merger litigation, and so forth? And so it's kind of helpful to break the world of merger review down into a number of stages.

So, you start out -- before the HSR, the agencies might not even be involved, but the parties are getting ready to do the merger. In principle, at least, they could hire economists to do UPP, GePP, merger simulation, what have you. In the real world, I think that's pretty rare. You usually come in sort of late in the game as economists, before the merger is filed but relatively late in the game. So, there may not be much role there unless the nature of lawyering and in-house counsels' offices and firms and investment banking were to change quite a bit.

Then the HSR is filed and the 30-day waiting period begins. So, the agency is going to be conducting interviews, and they will get responses to voluntary

request letters, which means the parties can produce what they want to produce, not what they have to produce. And in that window, you know, a large part of that may be devoted to just getting some information, and then you've got the remainder to processing the information. And practically, I think -- and I'll talk a little bit more about this -- it's going to be hard in that period to do anything too involved.

Then you have the second request stage, where the parties are complying. If the parties have chosen to do a rolling production, they produce data early, the range of things that you can do would expand dramatically. If they choose to just sort of wait until they have got everything and say, "Here you go, you have 30 days to decide to sue us or not," there is going to be a lot less that can be done, and that will sort of shift the favor between merger simulation and GePP, on the one hand, versus UPP on the other.

And then once the compliance happens, I think there's 30 days to issue a complaint or let the merger go through, and then things can move really, really fast. So, I went and checked the FTC's website, and the Whole Foods complaint came out in June 2007, Kevin Murphy's report was dated August 23rd, 2007. So, he had, like, two months to put it together from the time

at which he had all the information that was ever going to be available to him in his hands. I'm sure he was doing work before -- you know, in advance, but it's a relatively tight, compressed time frame.

And then if there is going to be a full trial on the merits, no matter who wins the preliminary injunction, then you're actually going to have a long time, and the range of things you can do is likely to be somewhat -- pretty broad.

So, I thought about this paper in the framework of competition among merger review tools, and hopefully competition will produce better results over time. The key advantage of UPP is that it's really intuitive, I think. When I read the formula or the description of GePP, generalized pricing pressure, it's actually pretty complicated, because you talk about holding the merging partners' response fixed but allowing all the nonmerging firms to adjust their pricing.

And I may be a little bit jaded by a recent anecdote that I'm going to share with you, but I tried to explain in writing to an attorney, a relatively smart attorney, that said when firms increase price, they face a trade-off. If they raise price, they will make more money on each unit they sell, but they'll sell fewer units, and a merger will weaken that latter incentive,

because some of the units that they would otherwise have sold -- have lost go to their merging partner, and those come back into their bank account via the other pocket postmerger.

And the attorney's response was, "That seems a little bit of a stretch. Can you simplify it a little bit?" Now, I want to be very fair to the attorney. The attorney understood it really, really well. What he was doing was looking forward to a district court judge, who has never tried an antitrust case, which is going to be the facts on the ground in many circumstances, and saying, you know, what's going to sound compelling and tie it to the facts of the case in the real world versus what's a bunch of egg-headed economists, you know, writing formulas and making models?

I said, "This is sort of the upper bound on intuition, and it doesn't get any simpler," and I think that that's probably true. So, there's a lot of intuition into it, but, of course, better intuition is only helpful and valuable in policy if it's correct intuition, and I think that's where there's a really important contribution made by the paper at hand.

In particular, an area for fruitful future research is to actually dig more deeply -- and this may require either simulation or ex post merger review, two

things that the authors themselves suggest -- to say, under what circumstances are you going to get the same answer when you take the simple approach versus the more rigorous approach versus when would you get the different answer and when would that lead you -- when would UPP possibly lead you astray?

And I have in mind almost something that you could perhaps guide the development of a checklist, similar to the coordinated effects checklist that says, you know, UPP is most reliable under these circumstances. When pass-through rates are high or low or something like that is going on, it's more apt to give wrong answers. So, you think you can improve UPP and be a complement, not necessarily a substitute.

There are -- I think I'll move quickly, because I'm almost out of time here.

I sort of went through the chronology of a merger. There's a time when data and time are both short, and you really can't do anything fancy, and then there's a time in most cases, some industries, like hospitals and airlines, there's public data and lots of agency experience with them, but in most cases, in the early stages, you can't do much fancy. In the later stages, you can do something fancy.

And I think the other point there is that at

that stage, there is not really competition between UPP and generalized pricing pressure. I think there's really competition between generalized pricing pressure and full-blown merger simulation models.

And so one thing we know about merger simulation models is when you look at what they predict versus what happened, they don't always do that well. And Craig Peters has a great paper on that, and I think, you know, one could actually revisit his analysis using GePP and try to determine whether the predictions would match any better -- would match more closely to what actually happened.

So, I think, again, you know, this is insightful, it helps sharpen thinking, and I look forward to more extensions from it.

(Applause.)

DR. DRANOVE: Sonia, did you want to just take a minute or two to react?

MS. JAFFE: No. I don't think I have any questions.

DR. DRANOVE: No? Okay.

Are there any questions?

(No response.)

DR. DRANOVE: This is a lively audience. Do we need to have incentives, immunity from FTC action for a

year if you ask the best question? I can offer that.

Our final paper in our early afternoon session will be presented by Thomas Jeitschko, from the Department of Justice. And, Thomas, I hope I didn't butcher your name. He's going to present Patent Pools and Product Development: Perfect Complements Revisited. The discussant will be Jay Pil Choi from Michigan State University.

DR. JEITSCHKO: Thank you very much.

Okay. Thank you very much for the invitation, and I hope you're still remaining to stay awake before the break one last time.

This is a paper written jointly with my co-author who's also here, and I for my part have to give the standard disclaimer, that these are my views and I'm not reflecting any views of the Justice Department. And I'm sure if you ask my co-author, she will also affirm that these are not the official positions of Towson University either, so...

So, what we're looking at here is we're revisiting a question of patent pools and looking at it in light of product development. Let me give you a little bit of background on this.

We started this paper a while ago, and so when I say "recent debate," you have to think of a slightly

larger time horizon here, maybe the last ten years or so, that we have become concerned about the patent thicket and sluggish innovation, the innovation rate being hindered, blocking patents, and all these notions.

In a lot of instances, this was tied to business methods or software, electronics oftentimes, and, in particular, also, biotech, were areas where this was noted to be a concern and a problem of stifled innovation.

This essentially led to what has now culminated again in IP reform, but there was a lot of discussion up to that, and one of the things that oftentimes was mentioned is the issue of cross-licensing or, in particular, what I want to focus on is the issue of patent pooling.

So, a patent pool is actually in an instance, if you have a collection of separate IP that might actually be relevant to produce a product, if the patents are all formed into a pool, then essentially these can be jointly licensed rather than people having to individually license across all individual relevant patents for their product.

This was essentially the idea that if you have to have access to a lot of patents in order to develop a product, then it would be incredibly costly to actually

try to gather and negotiate with all the different IP holders -- might even be very costly to identify them -- and it was in that context specifically that patent pools were thought to actually be a means or a tool to further innovation, foster innovation, in areas where you rely heavily on many different forms of IP.

Let me give you a little bit of a historical background, because I think as any antitrust person would notice immediately, you do have some thoughts a little bit. Well, if I'm pooling these things and now jointly licensing, are there concerns there? And, in fact, the first patent pool actually came about in the context also of stifled innovation. This was the famous sewing machine patent. There were conflicting patents. It wasn't clear exactly if they overlapped or not, and essentially, it was deemed it would be easier if you just form a pool and license them jointly.

The first time that there was a legal dispute over it specifically was at the beginning of the last century, at least here in the U.S., where there was a question of is this legal or can these upstream suppliers essentially -- these are upstream suppliers of IP -- can they simply form a monopoly over this IP or is there an antitrust concern there?

And at the time, there was a strong affirmation

that IP gives you these rights, and essentially the freedom of contract would trump any antitrust considerations in this context. So, basically, freedom of contract that is embodied or that is guaranteed to you through patent law and through the Constitution would trump any antitrust concerns.

Now, the first time when that was sort of revisited a little bit was not too long thereafter in the so-called bathroom trust cases, and in these instances -- so, this was about enamel, porcelain bathroom fixtures. There was also a trust there. In a lot of these instances, actually, it wasn't necessarily clear what was being argued here. Was it the behavior of the trust or was it directly tied to them being a pool?

That's also very much the case in the so-called Edison trust or the motion picture trust, that was also come about. There were conflicting or overlapping patents, they had formed a patent pool, but, again, a lot of the concern tied to the litigation in those cases was not actually the pool formation in and of itself, but some of the practices that the pool was exercising.

And in that instance, the pool itself was then found in violation of the Sherman Act, although, again, as I say, it's not clear that it would have been the

pool formation that was critical here; it was some of the practices that the trust engaged in beyond that.

Now, after that, the first case is essentially the Standard Oil Company. In a lot of cases after that, where we started to differentiate or the legal realm has started to differentiate that you really have to look at what constitutes the patents that are in the pool, what makes up the portfolio that's in it.

And, in particular, if someone has a patent that would essentially allow you -- would be a necessary input in any production, then essentially this would be termed a blocking patent. That means that the owner of that patent could prohibit anybody else from producing the downstream goods that require this input. And the concern was that you could have competing blocking patents. If they have sufficient overlap, you would essentially not be able to really produce anything downstream.

And the case law developed what were sometimes called competing patents or blocking patents or essentially, in the way we would talk about it as economists, would be we would differentiate between whether the patents were complements or substitutes. And, essentially, if patents are substitutes, that is, I need either one or the other, then a pool formation

would essentially just be forming a monopoly over this critical input, whereas you would have had competition beforehand. Otherwise, if the inputs are complements, then a pool formation should be viewed okay, because you need both of them anyway, so you might as well just license them jointly.

And we know we have exactly this distinction not just in IP, we have it in a lot of other instances. In any cases where you're looking at mergers of upstream firms in an industry, we oftentimes distinguish between whether these firms are producing complements in the supply chain or substitutes in the supply chain. And I think some of what I have to say might also address that a little bit.

So, what we've since sort of received as a conventional wisdom of how we make this distinction is if patents are perfectly complementary, so that you need both of them, then essentially, if the royalty rates for these patents are set independently, the problem that we have is an issue of double marginalization. In this case, it's a vertical form. We don't have monopoly rents stacked up upon them, but we have them simultaneously, two different sources and upstream.

So, this double marginalization, or which in this context, in particular, is then also sometimes

referred to as royalty stacking, reduces actually both producer and consumer surplus, or put another way, in a more stark formulation, if you actually allowed the upstream IP owners to pool their patents and you're perfectly happy with them even behaving as a monopolist, even allowing them to behave as a monopolist would increase consumer and producer surplus in the downstream market.

So, this harkens a little bit back to an original model by Cournot, who discusses this in the context of producing brass, where the two inputs are copper and zinc, and if you had an upstream provider of zinc who was a monopolist and an upstream provider of copper who's a monopolist, if those firms got together, you could actually increase total welfare.

So, the idea, then, is that whenever you have complementary inputs, if you can pool these together, even if you were concerned about the worst case scenario of them then licensing this with monopoly power, you still have an internalization of a pricing externality that you have otherwise, and you can increase total surplus.

Now, this model -- the insight is straightforward, and models of IP have actually used this in a lot of contexts. There are some things that

we felt were missing in the models of IP and downstream product development that follows, and so we wanted to revisit this question a little bit and look at a richer context and see if you really come to the same unqualified conclusion.

In particular, if you look at a lot of models in the theoretical literature of product development and IP and the role of IP, in a lot of instances, IP is essentially only viewed as some ingredient that you need, and that's also essentially what it was here. You need copper, you need zinc, and you can produce brass, in the Cournot setting. And essentially, that's all you really need to enter the downstream market, is you need access to these.

Then the models are then further built, so I think both Cournot and I think also in Shapiro's illustration, the downstream market is a monopolist who then produces the product in question and sells it, or in a lot of instances, you look at perfectly competitive downstream markets only to look at the incentives that happen on the pooling level at the upstream market.

Now, in contrast to this, one of the important things is in a lot of industries, IP is a necessary but far from sufficient input in your production, and so what we want to look at is, what happens if you consider

further development of products in the downstream industry and incorporate that in your model? And so what we're trying to get at is what we're going to call a more complete picture.

So, we want to develop a more comprehensive model in which we want to analyze three questions, then: We want to revisit the conventional wisdom in light of the fact if you imagine an industry that uses and really has to start working with the upstream IP. The second thing is we want to see what if pooling, in and of itself -- so, the process whether you actually access the IP independently across independent licenseholders or if you access it through a pool -- what if that actually has an impact on your development, on the way you develop and commercialize a product? And then to the degree that that can be the case, when might it negate the conventional wisdom or when might we revise our insights on the effects of pooling?

So, this is essentially our agenda, is we're going to build a slightly richer model where we focus, in particular, also on the question of downstream development and commercialization, and then ask the question of, what if the process of pooling actually affects that?

So, the model setup, a relatively easy

framework. We are going to look at two upstream patentholders, and we're going to assume the patents already exist. So, we're not interested in that aspect of the model. The patents are perfectly complementary. You need access to both of them in order to produce anything. And the patentholders have a choice. They can either form a pool or remain independent.

When we look at the downstream market, we are going to assume that these are differentiated firms in the downstream market who compete on prices, and so we have the standard model here, say Aviv's model. What we do add to it, though, is a development stage where essentially we say sort of the value of your own market depends on efforts that you have to put into research and development, and we postulate that this research and development effort might entail spillover. So, the size of your market or the desirability of your product is a function of your effort, but there's also spillovers in the R&D process, so it's positively tied to the efforts of your rival, and vice versa.

For simplicity, we are going to assume there are no costs in this model, other than the costs of providing effort in the research and development stage, and in the production stage, there are the licensing costs of IP. We are going to look at two instances. In

the case of fees, in some sense, we're going to think of this more as a benchmark. We know that the double marginalization problem comes about exactly when you look at royalties, not when you look at fees, and under fees, you have the incentives aligned perfectly between IP and downstream firms. So, we're thinking of it more as a benchmark.

It turns out, in most of these industries that we're looking at, royalties are the more relevant question, but we do want to think about this anyway, because we can see how the model works a little bit better. The other form is royalties on a per-unit basis, the output that is ultimately sold in the market.

The first thing that we find is that even if you enrich your model and you step away sort of from assuming that the downstream market is a monopoly itself or that the downstream market is a perfectly competitive market, even if you look at a differentiated goods downstream market where there's effort, it turns out that the conventional wisdom holds quite solid here. And that is that if you have royalty stacking as an issue, then even in these richer models, that can be overcome, and pooling is actually a good thing.

Now, the question we wanted to ask, though, is what might happen if, in fact, whether or not your pool

has an impact on the research and development stage? How could this affect it? Could the process of pooling, in and of itself, have adverse effects? And why might that be the case?

Well, essentially, in a lot of areas, we know that patents, in themselves, are characterized or -- patents aren't characterized, but knowledge, relevant knowledge is oftentimes captured by what is called tacit knowledge, and IP rights are therefore in some sense only an incomplete capture of the necessary IP that you want. So, this has been oftentimes discussed, that frequently, this is exactly what underlies collaboration between downstream developing firms and upstream IP holders, is the fact that just gaining access to the IP, in and of itself, is frequently not sufficient in order to really do something.

So, what happens then is implicitly, in using IP, where you also collaborate with upstream IP holders, you have a knowledge transfer that takes place, and that is essentially the movement of the tacit knowledge from the upstream IP to the downstream firm that then embodies this and can develop it and commercialize it. And here, the issue is that this process, in and of itself, there is reason to believe that it might very well be sensitive to whether the upstream firms, now the

IP holders, collaborate also with each other in the sale and delivery of the tacit knowledge or not.

To the degree that happens, I want to show a little bit of what can happen. So, in particular, we're going to think of the following: If you then think of a patent pool as an information-sharing institution as well as where the IP holders exchange additional information that they then jointly sell and package to all downstream firms, what you have is you might have spillover effects. That is, these firms are now operating with essentially the same team.

So, their R&D teams, their research and development teams now have the same input from the pool. So, that could increase spillovers. It can also diminish product differentiation in the downstream market. So, if they're on a more similar research trajectory, their final products might also be more similar.

So, this is essentially the model that we have in mind, is if the upstream IP is pooled, then in the research and development phase, you have larger spillovers in development, because essentially you have a research team that is accessing both of these developments. And once you have this, once you have greater spillovers and you have this, you might also

have more similar products in the downstream market.

So, looking at this structure, the first thing we want to look at is, how does this affect the downstream competition, the downstream development? The first thing is, what we know here is the amount of effort that you put into your research and development is a function of the spillovers, and to the degree that how spillovers affect it is a function of how closely your products are competing.

If you have close head-to-head competition, then essentially you're worried about your rival free-riding off of your research efforts. This is not a concern if you have highly differentiated products. And so that can flip essentially. It depends on -- how closely your products compete in the downstream market will depend on how your research efforts are affected by increases or decreases in spillovers. And that can also affect the market size.

So, essentially, what we have here is market size and research efforts, if the products are generally fairly close head-to-head competitors, are negatively affected by increases in spillovers. If you're worried about free-riding of a rival, it reduces your incentives of actually putting in an effort upstream.

Now, what that means is that this is a marginal

effect, and essentially what we have is sort of if you look at the space, there are no areas where essentially you have a lower total welfare, lower consumer surplus, and lower profits if you increase spillovers.

Now, if we look at what does the differentiation effects have, here we also -- this is unambiguous. That is, your efforts are going to be decreasing the more closely are your products. So, this we have -- independent of spillovers, you have this effect here.

Now, what this also means, this is on the margin, you also have areas now where closer competitors' head-to-head competition can actually be bad for consumers, because it reduces the effort that goes into research in the upstream market.

So, these were marginal effects, and what we want to point out is that essentially, in the downstream market, you can have adverse -- I only have two more minutes, so let me go through -- let me go through the main results, actually, because it turns out that this model involves a lot of different trade-offs that can take place within it, but one of the most important things might be we have a lot of consolations.

What I'm measuring here is essentially what's the degree that spillovers might be -- I'm sorry, that spillovers might be affected by pooling and the degree

at which product differentiation might be affected by pooling? We can see there's a lot of areas. If you start off with relatively homogenous products, then small differentiation effects will not affect your conventional wisdom, but otherwise, it might. And so you can have large combinations where pooling actually adversely affects your input.

I have here an example that, out of the interest of time, I am not going to go through, but we also can construct certain pathological areas, where you're worried about firms want to pool, even though it makes consumer surplus worse off. There are a lot of instances where firms want to pool, but it makes consumer surplus better off, so it doesn't matter. There are also instances where firms don't want to pool, but in those cases, consumers would also be worse off, so you are also not that concerned about it.

Let me maybe go to my conclusion at this point, though. So, what we have -- and the reason I'm happy to skip over these are because it's not clear to us entirely how pathological are these cases and what's the most likely one that you want to have, but I do want to conclude, however, that we think that pool formation can affect the downstream research/development/ commercialization stage, and we do think that it can do

so in an adverse way.

And, in fact, we think that this is particularly the case in industries where tacit knowledge matters a lot, and we think that that's also probably a reason why, in particular, in biotech, we don't find a lot of patent pools, even though they're advocated for them.

So, essentially, if patent pools make the products more closely similar to another and generate a lot of spillovers, then essentially this could mean that research efforts are going to be much more diminished compared to a nonpooling situation, even though in the model we assume that you have perfectly complementary upstream IPs.

There is some recent evidence, actually, that shows -- that looks at -- empirical work that shows that actually around pool formation, sometimes innovation actually slows down a lot, and that would be consistent with some of the concerns that we have here.

So, this was essentially our model about patent pools. Of course, I think it gives a little bit of thought also for other settings where we're worried about the combining of complementary versus substitutes, and in an upstream market, if the combination, in and of itself, alters anything about the downstream competition, you have to account for that effect as

well. So, the results might not be as strong as the conventional wisdom would have suggested otherwise.

Thanks.

(Applause.)

DR. CHOI: Okay. So, Thomas gave a nice presentation on an important topic. So, in this paper, he investigates private and social incentives to form patent pools in a richer framework, and then he challenges the conventional wisdom and derives some antitrust policy implications.

So, the conventional wisdom literature is that whether a patent pool would be procompetitive or anticompetitive would depend on the relationship between patents included in the patent pool. So, in particular, if the patents are complementary, then social -- actually, the patent pool will be procompetitive.

So, Thomas' paper actually looks at the following situation: So, in a sense, the patents are not complete in the sense that there should be more development effort to commercialize the -- finish the product. So, there is -- so, innovation in the literature -- I mean, in the paper is that there be a condition in the downstream firm, and then there be some spillover between two firms, and also, there would be some product differentiation in the final product

market, okay?

So, the first issue derives that even in this original framework, actually, the conventional wisdom will hold, unless there's no difference in the spillover rate and the product differentiation, okay? However, once patent pools induces higher spillover rate in the downstream stage and also product differentiation becomes smaller, then the conventional wisdom may be reversed, okay?

So, let me just comment first about another situation. Here, the fundamental assumption is the spillover rate increasing as a result of a patent pool, and also product differentiation becomes smaller, okay? So, basically, one important question, if you -- this is such a fundamental assumption, so what is the underlying mechanism for these changes? I think there should be more discussion about the paper, because everything hinges on this particular assumption, okay?

So, one justification given in the paper is that patent pool is a conduit for knowledge transfer, okay? So, if you go back to the paper -- I mean, to the diagram, so basically pooling is interaction between upstream stage, between firms k and l; however, spillover rate is at the downstream stage. So, the typical scenario given for spillover rate in the

literature is that maybe the scientists in these firms, maybe they go to conferences and they talk to each other or maybe the scientists may be publishing in some kind of a journal. So, that's how the spillover takes place. That's kind of the usual story.

But here, patent pool is at upstream stage, while spillover rate is in the downstream stage. So, the formation of pooling doesn't affect any mechanism we think about. So, basically, if there's some high spillover rate, that should have come through the upstream firms, okay? Otherwise, I mean, I don't see why there should be any changes, okay?

So, one story Thomas told was that actually patent pools -- okay, so these two firms have access to technology of both firms, but remember that here the technologies are complementary. So, even with independent firms, still, each firm will have access to both technologies. So, it is not clear to me why there should be higher spillover rate as a result of pool formation.

Maybe once the -- it might be because of information sharing at the upstream stage, there would be more knowledge transfer to the downstream firm, okay? So, that could also induce a higher spillover rate. But under the situation given in the paper, actually, after

a patent pool formation, there is no higher technology transfer from the upstream firm to downstream firm. So, there seemed to be a little bit of inconsistency. And I'm not saying that the assumptions are wrong, but I would like to see more discussion in the paper, okay? So, that's one.

Let me see. Also, the other thing is if there are some changes in the spillover rate, one question might be how important this is in the pool and what are the policy implications, okay? Maybe the test might be some empirical evidence, that would be great, but I doubt that there would be any empirical evidence. If there is no empirical evidence, there may be some anecdotal evidence or some inductive argument, okay, arguing for why there should be higher spillover rate. That would be highly desirable.

And then let me go to the model, okay? So, here, the paper analyzes the effect of a patent pool for fixed fee and the royalty rate cases, and in this paper, actually, these two scenarios are actually taken as kind of endogenous, okay? However, there's a lot of literature considering -- I mean, what would be the ultimate contractual form?

In other words, where the fixed fee will be chosen by the licensor, where royalty rate will be

chosen, these are all actually endogenous choices. Actually, I mean, the choices might also reflect informational constraints and also market competition conditions, okay?

So, for instance, here, the model -- I mean, the paper actually has more to do with the biotech industry, where actually there is some tacit knowledge. So, it is essential for the downstream firm and the upstream firm, they work together, okay?

Especially in this scenario, the ultimate contract or formula actually -- will be actually more like a royalty rate, the reason being that if you have a fixed fee rate, okay, so there's no reason for the upstream firm to put effort to make sure the downstream firm understands the terminology, okay?

So, actually, one of the results in the paper is that actually, the patent pool can be anticompetitive only when -- only when licensing would be fixed-fee case, okay? So, actually, in the royalty rate case, actually, that kind of concern does not arise.

So, if this is the case, in other words, if technology transfer of tacit knowledge is more important, then this theory predicts that there should be royalty rate. So, that also implies that there should not be that much concern in the paper.

Okay. So, let me just keep on moving. One line of research actually I'm engaged in, okay, which is there may be another mechanism where patent pooling might be actually anticompetitive, okay? One is kind of like the idea, patent pool might be a mechanism to harbor a weak patent. So, we can think about probabilistic of patents. So, let's assume this story. There are procomplementary patents, A and B, so there is -- and independent firm C, actually, based on patents A and B, develop further technology, okay?

So, let's say I file a paper on probabilities of validity in the court, okay? Let me assume that -- let's say two patents are very weak. So, basically, the patentholder, they would like to eliminate the incentive to litigate, okay?

So, in that case, firm C, the incentive to litigate would be -- this would be applied really in the patent, okay? They can save a fee of fA, and the L is, let's say litigation cost, okay? So, as long as this condition (inaudible) them, there would be no incentive to litigate the patent from C.

So, then, the highest licensing fee that firm A can charge will be given by this number, okay, and this would be the highest fee that can be charged by firm B, because that is kind of a limited royalty fee. That

would eliminate any incentive to litigate.

Okay. Now, let's consider patent pool and the data for some kind of joint defense, okay? So, here, remember that there are -- these two patents are complementary, so what that means is that if there's a patent pool, then the only reason -- the only way firm C can avoid paying a royalty rate would be that the firm invalidate the first patents, okay?

So, then, a patent pool challenge will be -- let's say patent pool challenge of f, and as long as this is the product that the firm C will -- invented with the first patents, okay? So, this is the conditions that will eliminate the incentive to litigate. So, basically, this is the fee -- the patent firm can charge, why this would be the fee, but coming from royalty, coming from independent licenses.

So, as you can see here, with patent pool actually you can charge a higher price. So, this would actually be reverse of the Shapiro work. So, in other words, another mechanism you can think about might be introducing probability patents. Then a patent pool can be used as kind of a mechanism to harbor weak patents. So, that could be another mechanism in which a patent pool might be anticompetitive, even though two patents are roughly complementary.

So, let me just -- a summary. So, this paper provides a new perspective on patent pools in a richer framework, and those are very clean and elegant characterizations. So, one concern is the endogeneity of the licensing contractor focus of this.

(Applause.)

DR. DRANOVE: Thomas, wherever you are, a quick response?

DR. JEITSCHKO: (Off mic.) Thank you very much for those comments. That's helpful. I think one of the -- one of the points I'd like to make, because of the -- I'm not sure whether it came over exactly right. So, we do have instances where under royalties, also, you have lower welfare in total, and so we had an example where we could characterize where royalties aren't a problem, but we also have examples where they are a problem.

DR. DRANOVE: Okay, thanks.

Experience has taught me that I know you all are bubbling with enthusiasm to ask questions for both the speaker and the discussant, but, alas, we are running over time. So, you'll have to catch up with them during the break. Thank you all.

(Recess.)

DR. ADAMS: Moderating this session is one of

our scientific committee members, Nancy Rose.

(Applause.)

PAPER SESSION THREE:

CONSUMER DECISION-MAKING AND SELLER INCENTIVES

DR. ROSE: All right. So, we've got a set of three extremely interesting papers that I am confident will rejuvenate the audience, bringing you back from that break, I hope you all got a chance to choose your favorite form of caffeine, if you need that, or a little sustenance, but a very interesting selection of papers on consumers' decision-making and seller incentives. And we'll kick off with Matthew Osborne of the Bureau of Economic Analysis with a paper on cellular service demand that will be discussed by Eugenio Miravete from the University of Texas at Austin.

DR. OSBORNE: All right, very good. I apologize in advance if I cough a bit during this presentation. Unfortunately, I got horrendously ill before this, and I am over the illness, but my voice has not quite recovered.

So, anyway, this is joint work with Michael Grubb at Sloan. I am going to start out with everyone's favorite slide, obviously, the disclaimer. So, these are our own views and not the views of the Department of Commerce or the Bureau of Economic Analysis.

So, what do we do in this paper? Okay, so what we're going to do is we're going to estimate a model of

tariff and usage choice in some very interesting and detailed cellular phone billing data, okay? And this model has two features, which I think are very interesting.

So, the first feature, which I think is novel, is that we have a nice way of incorporating consumers having ex ante uncertainty about marginal prices for phone calls. And the idea behind this is very simple. When you're on your cell phone and you're making a call, you know, you may not know if you're going to be over your minutes or not, right? So, the call may be zero cents per minute or you may be paying an overage fee, okay?

So, in our model, we recognize, you know, it's hard for people on cell phone plans to track their usage over time. So, what they're going to do in our model is they're going to set up what we call an optimal threshold rule, where they'll set up basically a rule of thumb where they'll take calls that are very important to them and then they will reject calls that don't seem so important, okay?

The other thing that we do in the paper, which I think is pretty interesting, is we're going to identify what we call bias beliefs, okay? And so the idea behind this is in our data, we see everybody's bills and we see

everybody's usage, okay? So, we can actually identify the true distribution of tastes, because we see that in usage.

We can also identify what people's prior beliefs about their usage are going to be, and the intuition behind that is that people are going to sort into plans depending on how much they think they're going to use. So, if you think you're going to be a low user, you are going to sort into a small plan. If you are going to be a high user, you are going to sort into a big cell phone plan.

We see people over time in our data, and so we observe people picking, say, small plans, and then, over time, switching into bigger plans if they are actually big users. And so we are going to be able to estimate how quickly people learned about their true taste.

And what we call biases are differences between people's beliefs about their usage and their actual usage, okay? And we're going to talk about what -- we're going to quantify what these differences are.

And then we're going to do two counterfactuals, which I think are going to be pretty relevant and interesting. So, the first counterfactual is going to be, you know, how costly are consumer biases? So, basically, if we were able to take away people's biases,

how much better off would they be?

And the second counterfactual we're going to talk about or that we're going to quantify is some bill shock regulation that the FCC I guess is now going to implement in 2013, and the idea behind what this bill shock regulation is going to be is that the FCC has agreed with cell phone companies that they're going to send you a text message when you get over your minutes. So, you're going to know -- now you're not going to be uncertain. You're going to know, hey, I'm going to be charged 45 cents a minute for the calls that I'm making.

And what we're going to show in the paper is that if firms can endogenously adjust their prices in response to this regulation, consumers can actually be worse off under the regulation, and that's also going to tie heavily into the amount by which consumers are biased, all right?

So, before getting into the model, let's talk data. So, what kind of data do we have? Well, we have individual-level cellular billing data for two years for all the students who subscribe to cell phone service through a major U.S. university. We also have pricing data for all the cell phone carriers that operated within the area of the university during that period, and we used that in our supply side estimation.

Now, to give you an idea about what the data actually looked like, so what the university did was it contracted with a cell phone carrier, and it offered a set of plans that were somewhat different than the plans that you could get through the public, okay? So, for instance, here's a sample of what the four most important plans looked like.

So, the first plan, which we call Plan 0, is basically a two-part tariff. You pay $14.99 a month, you got no included minutes, and you paid 11 cents per minute for all your calls. And normally, these sorts of plans are not available.

What we call Plans 1, 2, and 3 are more traditional three-part tariffs, okay? So, these are plans where you -- so, Plan 1, for instance, is a plan where you pay $34.99 a month, the first 380 minutes are free, and then any minutes you use after 380 minutes, you pay 45 cents a minute. And these went up to Plan 3, which was a $55-a-month plan, where you pay 875 minutes -- or you get 875 free minutes, and then you pay 40 cents a minute for the rest of them, all right?

So, let me talk through the model. I'm going to give you a very simplified version of the model from what's in the paper. The paper's model has 30 parameters in it, and it's really complicated, and this

one's going to be a lot simpler.

So, the basic way that the model works is that at the beginning of the month, consumers are going to choose a plan, j, okay? And then they're going to decide on a calling threshold, which we're going to call v-star, okay? And this calling threshold is going to be based on their beliefs about what their taste for calls during the month is going to be. And we call that theta. They don't know what their taste is at the beginning of the month. They only see that at the end of the month when they get their bill.

Now, the way you can interpret this taste for calls, theta, is it's basically a measure of all the calling opportunities that arise during the month. So, throughout the month, theta calling opportunities arise, and theta is the total number of calls you could make if you didn't restrict yourself at all from making calls. So, it could be, like, 2000 minutes or 1000 minutes or something like that.

And what's going to happen is throughout the month, we're going to assume people can't track how many calls they've made. They're just going to see the value of a call, and they're going to reject it if its value is below v-star, and they're going to accept it if it's above v-star, okay? So, what that means is people are

going to make some fraction of all the calls they could make, okay? And that fraction is going to be called q-hat of v-star, or we're going to it can q-hat of v-star.

So, what this means is at the end of the month, a consumer's usage is going to be q, which is going to be theta, the total number of calls they could have made, times q-hat of v-star, which is the fraction of calls that they actually accepted, okay?

Now, in our model specification, q-hat of v-star is going to be one over one plus beta times v-star, okay? So, you can see that this is a decreasing function in v-star. As your threshold v-star goes up, the q-hat goes down and you make less calls.

Now, where does this come from? The way that we get this q-hat of v-star is we specify a utility function for calls, okay? So, we specify a value of minutes which looks like this, and then we recognize that since v-star is the value of the marginal call, it has to be the case that at theta q-hat of v-star, v-star is equal to the marginal value of theta q-hat of v-star. So, v-star has to be equal to the derivative of the utility function here.

And so basically we know what this derivative is, we can solve it, and so we can back out q-hat of

v-star from there. And so that's where all -- that's how everything sort of fits together, all right?

Now, let's talk about biases. So, people have beliefs about their taste shock, theta. Now, theta's measured in minutes, so it's got to be positive. So, we assume that theta follows a censored normal distribution, all right? So, what that means is that there's some underlying theta tilda. Theta is going to be zero if theta tilda is less than zero, and it's positive if theta tilda is greater than zero.

The underlying latent taste shock, theta, is going to have two pieces. It's going to have an individual fixed effect, which we call your true type for calls, and that's what people are learning about. We're going to assume they don't necessarily know that. And it's going to have an idiosyncratic error, epsilon-it, all right? And these are both normally distributed.

So, in reality, the theta tilda is going to be normal with mean mu i and variant sigma tilda epsilon. Now, we're going to assume that people's beliefs are that it's normal with mean mu i and variant sigma tilda epsilon. So, we're going to allow their belief about this variance to differ from the actual variance by some factor, delta.

Now, if delta's one, then that means those variances will coincide, and we would say that people are rational in that case. If delta epsilon is less than one, then we call that projection bias, and what that means is people underestimate the volatility in their month-to-month taste variation, okay?

There's another type of bias that I'm going to focus on, which I call overconfidence, and it's similar because it ties into what consumers believe about their uncertainty about their true type mu i, okay?

So, in the model, every consumer is initially assigned a prediction of their true type or a belief about their true type, which I'm going to call mu tilda i. That's going to be drawn from some population normal distribution. And over time, people are going to update this mu tilda i 1, okay?

So, in period one, you get my tilda i 1. At the end of the period, you see your theta, and then so you update your mu tilda i 1 using Bayes Rule, and you do that period by period.

Now, at the beginning, when you first sign up, you know, if a consumer was sort of fully rational, their precision about the -- about the mu i would just be the conditional variance of mu i, conditional on the mu tilda i 1, where that's taken over the population

distribution of these things.

We're going to assume that people's prior beliefs of mu i, given their initial information set, are going to be mu -- mu tilda i 1, and the variance is going to be sigma tilda i 1, okay, where, again, sigma tilda i 1 is going to differ from sigma mu by some factor delta mu, okay?

So, if delta mu is equal to one, then we would say people are rational. If delta i is less than one, we're going to call that overconfidence, and basically what that means is that people underestimate the uncertainty about their own type. So, they're going to be too sure about what their -- about what their true type is.

So, I'm just going to give you -- I'm going to give you a brief overview of what the estimates are or what the important estimates we think are, and then I'm going to talk about the counterfactuals, and then conclude.

So, the price coefficient that we estimate is 3.4. That's the coefficient in the q-hat. And basically, what this means is that people are price-sensitive. So, if you increase the price of a call from zero to 11 cents per minute, people reduce their calls by about 30 percent, okay?

We find people are overconfident. So, if you look at sigma mu, the true variance of the mu, it's -- or the true standard deviation is 107 minutes. People's beliefs are 14 minutes. So, they're highly overconfident. And what this means is that people are going to be too sure about their types and they're going to sort into plans that are too risky. So, if you think you're going to be a 100-minute user, you're going to be really sure about it and you're going to sort into a plan that's too small for you.

We find the same thing with the idiosyncratic error, epsilon. So, basically, we estimate the sigma epsilon to be 169 minutes, but people actually believe this error is 91 minutes. So, they're going to display projection bias, and they're going to sort into plans that are -- and that's also going to cause them to sort into plans that are too risky.

There's other types of biases that we allow for in the model, and I don't have time to talk about them all here. One that I'll just mention is that we allow the distributions of the mu tildas and the mu's to differ, and we find that the means are systematically different, okay? So, the mean of the mu naught is 107 minutes. The mu tilda naught is minus 25.

Now, people don't believe that their usage is

negative. Remember, the tastes are censored normal. So, when you take the variance into account, people initially will believe that they're going to use about 30 minutes, okay, when their actual initial usage is going to be about 110 minutes. So, they underestimate their initial usage by about 80 minutes. And what this means is that they're going to sort into plans that are too small, on average, okay?

And just as a comment, we assume in our model people are risk-neutral. If you believe that people are risk-averse, then basically to rationalize what we would see in the data, you would need even more overconfidence and projection bias. So, in that case, you'd want to interpret these results as lower bounds.

Now, what do the overconfidence and projection bias mean? Well, the one way that you -- the way that you tell them apart in the data is that they affect the rate of learning, and the way that they affect the rate of learning can be seen as follows in this little simulation study that I've done.

So, we find that the overconfidence is a lot stronger than the projection bias, and so what that means is that people place way too much weight on their priors relative to the signals they get on their bills. So, they update their beliefs way too slowly.

So, what this means is -- I'm showing you here. Take a guy whose actual type is 107 minutes who starts out believing he's a minus 25-minute guy, okay? After about a year, he is going to have updated his belief to be a little bit over zero, okay?

Now, if we take away the biases, which is what's this line here, you can see he updates his beliefs much faster. He updates them up to about 80 minutes, okay? So, the projection bias and overconfidence are slowing down the learning a lot.

The biases that we estimate also have a big effect on plan shares.

Okay, I should be able to finish in five minutes. That shouldn't be a problem.

So, these are the plan shares at the estimates. If you take away overconfidence and projection bias, what happens is people sort away from risky plans, like Plan 1, into bigger plans, which is what we would expect to happen. If you take away all the biases, like the aggregate mean bias, people sort into bigger plans.

One anomaly you might notice is that the share of Plan 0 goes up, and that happens because of a bias we include about out-of-network shares, which I don't have time to talk about, but if you ask me after the talk, I can tell you about it.

The welfare effects of the biases are significant. So, in the data, the average bill is about $41, and the overage probability is about 20 percent. If you take away overconfidence and projection bias, the average bill goes down to about $38, and you can see the overage probability drops down by about -- by 10 percent, okay?

So, what's going on here is people are paying more in their monthly fees because they sort into bigger plans. Their overage fees, though, drop significantly, because people cut back their usage a lot or they make a lot less overages because they're in bigger plans. So, their overall bill drops by about $3, okay?

So, the firms -- what's going to happen here is firm profits are going to drop, and here I'm using profits and revenues as synonyms. I'm assuming, as a first-order approximation, that firm marginal costs are zero, consumer welfare goes up, and total welfare goes down. The reason total welfare goes down is because usage goes down.

Just as a side point, the reason -- one thing I didn't explain was the reason usage goes down is because when those variances go up, the v-stars go up. People recognize that the probability of an overage is higher, so they become more choosey about what calls they accept

and reject. So, they actually make less calls.

And since in our model total welfare is just going to be an increasing function of total usage, if total usage goes down, total welfare has to go down as well, all right? If you take away all the biases, then obviously consumer welfare gets even better.

If you look, these numbers are measured per person, per year. So, in total, people are better by $50 per person, per year, all right?

Just as a quick comment, if you look at the public plans, the results get even larger, and the reason for that is for people on -- for the people on the public, you know, they didn't have access to stuff like Plan 0. They could only sort into plans that were much more risky, like Plan 1, okay? And so the effect gets even larger.

So, let me conclude by talking about bill-shock regulation. So, the way that we simulate bill-shock regulation is we assume that there are three firms that each offer two different plans, okay? And what we're going to do is we're going to assume those firms are symmetric and we're going to solve for the symmetric equilibrium.

So, what we find in equilibrium is at the estimates, the firm offers two plans, okay? So, one

plan it offers is a two-part tariff, which is $30 a month and 50 cents a minute and no included minutes, and the other is a three-part tariff, which is $60 and offers about 300 included minutes.

Now, if you don't allow firms to change prices, what does bill-shock do? Well, in our model, what bill-shock is going to do is people are going to be following along their v-star threshold rule until they run out of minutes, and when they run out of minutes, they are going to realize, "Oh, I'm getting charged 45 cents a minute, so I'm going to raise my v-star to 45 cents, and I'm going to cut back my calls a lot."

So, what happens is people tend not to switch plans, but they cut back their calls a lot, and so firm profits go down. Because they cut back their calls a lot, total welfare also goes down. Consumer welfare goes up, because people are paying less. But things change when we allow firms to adjust their prices endogenously.

So, if you allow firms to adjust their prices endogenously in response to the bill-shock regulation, here's what they do. So, first, they leave this plan alone, okay? On this plan, people are going to be more sensitive to the cue, right, because of the bill-shock regulation, so they have to raise it. And to compensate

for that, what the firm is going to do is it's going to raise the fixed fee of this plan. So, basically what the firm is doing is if you look at the margins, it's trying to keep the margins across the plans at about $80 a person, per plan, okay?

Now, what happens is usage goes down, so total welfare goes down, but the firm is raising prices and trying to keep its profits the same. So, its profits, they actually go up just a tiny bit a little counterfactual, but because profit -- the firm is raising its prices and total welfare is going down, consumers are sort of stuck, because they're the residual claimants for total welfare. So, therefore, their consumer welfare has to go down. So, in this case, consumers actually get hurt by the regulation.

And just to conclude, here's what happens -- here's what happens when you take away the biases. So, when you take away the biases, it turns out we're finding that the optimal menus of tariffs are when there's no overconfidence and projection bias, you have two-part tariffs or you have a two-part tariff and a flat-rate plan; when there's no biases, you just have two flat-rate plans.

I don't have time to get into what's going on here, but basically this result is consistent with

Michael Grubb's AER paper where he shows that without biases, it's pretty hard to generate three-part tariffs, okay? And basically here, I mean, if there is no three-part tariffs, there is no skill for bill-shock regulation, so it's not going to do anything.

Okay. So, to conclude, we've estimated a model of tariff and usage choice using cellular phone billing data. We find overconfidence and projection bias. The biases are significant, and they have a big impact on consumer welfare. Bill-shock regulation helps consumers if prices don't vary, but it can hurt consumers if prices vary.

All right, and I will turn it over to the discussant.

(Applause.)

DR. MIRAVETE: Thank you.

Well, the social planner was certainly biased if they thought that I could discuss a 90-page long paper, including the appendices, in seven minutes, but I'll try my best.

Okay. So, it's a very ambitious paper. I like it a lot. The comment on nonlinear pricing not being dead, well, I sort of got that when I was in the market initially, when people -- everybody was, yes, estimating demand with differentiated products and things like

that. But, well, since then, Mark has been pushing for theoretical work, and myself and some other people are doing things on nonlinear pricing and glad that, you know, you guys continue doing this.

It's an outstanding data set, very sophisticated structural approach, and overall, there is a very interesting policy question. I think it points out to another example of unintended consequences. It's a very long and winding paper, and that's perhaps the -- what I guess in the long run needs to be improved in terms of the presentation, but it shows -- and I want to emphasize this -- that it shows a very, very good taste in how the modeling is made, how the data is analyzed, and how the features of the data are incorporated into the structural modeling.

The authors have the good taste of citing all my papers on nonlinear pricing, and that's also a plus. Anyway, at least I didn't have to wait 55 years, like Aviv was mentioning this morning for Gorman, to get -- anyway...

So, let me go and try to make some remarks. So, there are some issues in terms of how the paper is presented. Well, there are biases. We know people are making a choice of the plan first, and then they decide when to consume, and so on, but this is actually

formally defined on page 28, which I found it difficult -- there were many concepts into the -- discussed at the beginning of the paper. So, for instance, overconfidence is formally defined on page 29, and that makes less effective the presentation at the beginning.

Whether it is inattention or not, it's unclear, you know, it's very costly. Somebody could actually be making phone calls and have, I guess, measuring on the watch for how long you've been talking or how many calls you make. That's very difficult to do, just a matter of cost. It's not a question whether the individuals are inattentive or not.

The data is very rich. I wish I had it years ago. What else? As I said -- and this I really liked on the paper, the very sensible choices, exploiting features of the data, like, for instance, how there is a change in demand going at 9:00 in the -- p.m., and that's when you change the regime from peak to offpeak, and that helps to identify the price effects. And the paper goes into very -- a lot of detail in exploiting the data to make, in a sense, good decisions.

There's still issues about untested statistical assumptions on distribution and so on, and I don't know whether this would be relevant for the results. The

model is heavily parameterized. That's one issue that I think that the authors need to look at.

What else? Yeah, so I need to give a measure of this, you know, very useful analysis, ten pages. Discussion and authentication, I think it's one of the most complete discussions and authentications that I've seen on any paper, and it gets to a point where you may lose sight of what's going on in the paper, okay?

It's so detailed that later on -- and I really like the counterfactual, but you actually -- by comparison, it goes very quick on how you build the optimal tariffs and so on. I think that's actually very, very interesting for practical purposes. I mean, what are we doing? I mean, we want to look at unintended consequences. I mean, how do we build these optimal tariffs? But, again, by comparison, I think it stands out.

Easy fixes. The introduction -- maybe I can actually make it in seven minutes, I don't know how -- two, okay. A little overreaching. You're just trying to stop every potential criticism that you can receive of the paper. I mean, you are just doing what you are doing, and you are doing it fine. I don't know, it's a little bit too much.

Cellular service is not new. It's not a new

good in 2002. I have papers on the early nineteen eighty -- '86 or '84, I think, and I really don't think you need it. You just get into a new regime, a new tariff plan, and you need to learn the features of the tariff. So, that's already an experienced good.

And you may need a few months to figure out whether your consumption pattern is the right one given the tariff that you have chosen. So, it's not about learning on cellular usage.

You talk about, at some point, whether two different alternatives -- whether the distinctions of the (inaudible) are pedantic or not pedantic. You are discussing two potential structural models, and perhaps the example is when you go and discuss about risk aversion. On average, I mean, we are talking cell phones. Again, it's a -- in general, this represents a very low percentage of the total expenditure in the real world of income.

And it's, you know, a reasonable assumption to assume that individuals are risk-averse -- I'm sorry, risk-neutral. But, in fact, there is no way -- I mean, you give the alternative interpretation, well, we assume that we are risk -- that the individuals are risk-averse, the -- what we read today as overconfidence will be read as risk-lowering behavior.

Simply, you don't have the data to distinguish one structural model over the other. So, yes, pick one, and I think it's perfectly reasonable to assume that you have risk-neutral agent, and go on.

Consideration sets. So, you ignored the case of individuals who decide to stay in their own plans, and I miss that you use your microdata to figure out whether the individuals -- how many of those individuals would be better off by staying in their own plans.

And I hate citing myself here, but there is evidence in a different environment that individuals -- actually, Medicare -- when they don't switch, that's not proof of inertia. In many cases, they don't switch plans because they are actually in the best plan possible or close to the best plan possible. And you -- I don't know how much that would complicate the estimation, if you just look at those individuals who are better off by staying in the plan that they are currently subscribed to.

I think it's -- I understand the ex -- the two exercises that you are doing, but I think it's -- from a policy perspective, it's much more interesting, the one on the bill-shock, because you can just -- it's very difficult to make the bias expectations to go away.

The last one. So, I have the bad habit -- I

don't like standard errors. I wish everybody reports the statistics. Your statistics in Table 7, I believe, there is one that is not -- there is one parameter that is not significant. If you compute this statistic, it ranges in the thousands. The highest one, I believe, is over 4000.

So, that's a hint that maybe you have an identification problem, and, you know, essentially all -- the parameters are fixed variables there. So, I don't know which one -- I mean, I don't -- I am not arguing against your identification assumptions. I think they are all perfectly right. The only problem, maybe you don't have enough variation in the data to pin down so many parameters, and that's something to look at carefully.

That's all. So, that's it.

(Applause.)

DR. ROSE: So, we're I think a little bit behind time, but is there one or two questions?

AUDIENCE SPEAKER: Thank you for a very interesting paper. I'm Eileen Rule. I'm from the Federal Communications Commission, from the Consumer Bureau, and one of -- there are a couple of motivations behind the bill-shock work, and I wondered if I could ask you to address whether -- a second one, which is not

in your paper, I think, and that is the frustration that people experience when they have a dispute over their bill and then complain with their carrier and then complain with their regulator and then complain with us.

Just over the break, there was someone who told me about a two-hour conversation with Sprint. So, you know, that kind of -- that has an impact on consumer welfare as well, and I was wondering if there was a way for you to incorporate that.

DR. OSBORNE: Yeah. You potentially could, although I guess we would have to think about it a little bit. I mean, there is some -- so, you actually do see a few instances of this in the data, like there's one example I can think of where somebody started using roaming, and I think they didn't realize they were using roaming, and they got a huge bill. That's certainly a salient point that we can -- you know, that we can think about.

It's not something that we've really addressed, but, yeah -- no, I mean, I guess the welfare effects of a few people sort of potentially getting really screwed and having to pay $20,000 could swamp other stuff.

AUDIENCE SPEAKER: It's not just a few people, though. I think there's very few people in this room who have not had a frustrating experience with a cell

phone provider.

DR. OSBORNE: That's true. Yeah.

AUDIENCE SPEAKER: So, it's quite pervasive.

DR. OSBORNE: Okay, okay.

AUDIENCE SPEAKER: Yes.

DR. OSBORNE: Okay. Bo, we didn't think about billing errors, but that's a good idea, so...

DR. ROSE: I was going to call time, but if Joe Farrell's got a question, I'm not going to be the one to stand between him and his questions.

DR. FARRELL: You know me so well. Thank you. I'll just be relatively quick.

In the bill-shock analysis, it seemed as if you were holding the marginal price above the usage allowance fixed at 50 cents a minute.

DR. OSBORNE: Oh, we put a --

DR. FARRELL: I would have thought that one of the likely effects of better disclosure was to put more pressure on the carriers to lower that, in which case that could seriously affect your welfare calculations.

DR. OSBORNE: Yes. So, we explain in the paper, we actually have to put an upper bound on that price, because what we find is that the rational inatt -- or the inattention, combined with the biases, causes the firm to raise that price to levels that are

very, very high. So, what it actually wants to do is put that price high and put the fixed fee very low. And we don't think that that's entirely realistic.

I mean, in the population, at least, there's going to be some people who -- you know, who have been using cell phones for a while and who are going to be less biased. There are going to be some people who are going to be larger users, and it's going to be harder for the firm to do that sort of thing. So, we put that bound in there so we don't -- you know, as you know, I mean, demand analysis has to be local, and we don't want to make predictions that are sort of way out of line.

So, that's your answer on that.

DR. ROSE: Okay. I am going to call time so that we don't run too far behind on the others.

All right. So, we started with a paper on cell phones, which everyone in this room is presumably familiar with and eagerly awaited the results of that paper. We're moving on to another favorite market of economists, which is eBay. Henry Schneider from Cornell is going to be talking about his work on bidder preferences, and that work will be discussed by John Asker.

DR. SCHNEIDER: Okay. So, I'm very happy to be here and to participate in this conference. So, my

paper is about auctions, it's about eBay, and we'll be looking at nonstandard behavior.

And eBay, just to be clear, there are two categories of nonstandard behavior that we're really thinking about. One is making mistakes, bounded rationality kind of stuff, and one is more along the lines of rational behavior with nonstandard preferences. So, really, we're going to lump them together, and we'll have separate types of nonstandard behavior that we'll sort of break out in individual instances.

My co-author on this paper is Joe Podwol, who's a former Ph.D. student of mine in the Econ Department at Cornell. Anyway, so one other caveat before I really get started. This is not -- you know, it may at times feel like an antibehavioral kind of paper, and it's certainly not meant at all to be that. You know, so we will show a lot of null results, probably more null results than you usually see in papers, but that's not saying -- you know, it's null results with respect to tests of nonstandard behavior, but we're not saying that people actually adhere to sort of what's typically assumed as standard behavior, only that we can't rule it out, really, using sort of previously used tests, okay?

Okay. So, to get started, let me just briefly mention the laboratory work that's been done. So,

there's a fair bit of work, dating back 20, 25 years, maybe, that finds that especially in second-price auctions, we see a fair bit of overbidding -- not a fair bit, but a majority of inexperienced bidders come into the lab, and they tend to overbid, and by quite a bit. And this behavior tends to be pretty persistent.

So, more recently, there have been other lab studies that actually try and see whether bidders learn to avoid this sort of "mistakes" or nonstandard behavior, and generally, they find that there is circumstances in which people can learn to avoid overbidding. And so there is hope, actually, based on some of these papers. And so that's actually where we're going to come in. We're going to really focus on a real-world auction setting.

eBay, obviously, is a natural place to look, and we're going to conjecture up front that, you know, most bidders on eBay are pretty experienced. Even -- you know, the numbers we found, even the 25th percentile of bidder experience on eBay is something like dozens of auctions. So, if there's hope for learning or an experience effect to make these behaviors go away, we might expect to see it here, okay?

So, we're obviously not also the first paper to look at this issue. There's, you know, ten -- seven,

eight, nine, ten, a dozen papers that look at these issues, so we're going to add to this existing literature. We're going to argue that, you know, we're basically taking tests that have previously been used and try and, you know, maybe get somewhat better identification, I think is our comparative advantage. We do an experiment, among other things, to try and get cleaner test results. And so we're going to find results that are at odds with a sort of vast majority of this literature.

So, you know, we really -- you know, the version of this paper as it currently stands is very different than what we initially set out to test, actually. This was sort of a bread and butter IO paper at first, and then we found stuff that was sort of inconsistent with previous work, and so really the focus changed to really see, you know, why exactly we're getting different results from previous papers and to really say something hopefully meaningful about sort of the presence of nonstandard behaviors versus sort of more traditional behavior in this important market, okay?

Okay. Okay. So, as we all know, so eBay is well known as an auction marketplace, but -- so, by 2010 or 2011, there's actually a large fixed-price component of eBay. So, there are these things called buy-it-nows.

So you type in, you know, Batman Begins as a DVD, you want to buy it on eBay, and up comes all your auctions, but also up comes a lot of your fixed-price listings. And so, you know, at least in theory, you see all of them together, both fixed-priceds and auctions on one site. That's going to be important for some of our stuff.

We're going to focus on movie DVDs, primarily new movie DVDs, so we have a new product. We don't have to worry so much about unobservables in terms of quality and some other issues, although we will sort of broaden the focus a bit by the end of the paper.

So, let me tell you a little bit about some of the theories we're thinking about or maybe behaviors is probably a better way to put it. So, auction fever is one that, you know, I'm sure we've all thought about a little bit, at least. There's quasi-endowment effect, which has been proposed. It's sort of a variant of the endowment effect, and basically the idea is that -- a conjecture that bidders get some attachment to the item even before they win it, and so you get the traditional effects with just -- even without owning the item.

There is an escalation of commitment effect, which is sort of like a sunk cost fallacy story where, you know, you put all this time and effort into bidding,

and so that might lead you to overbid because of the effort that you've put insofar.

And then these two last ones are sort of in the spirit of sort of making mistakes kind of behavioral issues. So, nonrational herding says that bidders are herding into auctions with mistakes in their beliefs about what, you know, previous bidders say about the, say, the unobserved quality of the auctions.

So, really, Simonsohn and Ariely conjecture that you get a lot -- so, they provide empirical results that you get a lot of folks herding into auctions that already have a lot of bidding activity, and that extra bidding activity is actually due to lower starting prices, and that has actually very little to do with any kind of unobserved quality. So, that's a sort of mistake in the inference about the unobserved quality that bidders are making.

And then following it, irrational limited attention is probably the most high-profile paper that we're sort of looking at here, but that's the idea that bidders are ignoring these fixed-price options on eBay and bidding up the auction prices and in some kind of systematic fashion object to form these fixed-price options, okay?

So, generally speaking, there's two sets of

tests that people use to look at these things or at least two sets. We are going to focus on two sets. The first is looking at starting price effects, and I'll tell you a little bit more about that in a second, but essentially, the idea is -- this first test allows us to look at the first four of these behaviors and not at the irrational limited attention, okay?

So, the idea is basically you want to look at an auction that has a lot of early bidding activity, and the idea is that that's going to lead to -- that early bidding activity itself is a trigger for future bidding activity, okay? So, probably the easiest one to see is with nonrational herding, the idea is that, you know, you have a low starting price, a lot of people come into the auction early on, and that itself is going to attract people into the auction. They're making sort of biased inferences about quality and so on.

So, anyway, so that's sort of, you know, the test that's sort of been proposed. Obviously, there is some serious endogeneity problems with regressing, say, you know, past bidding activity on future bidding activity for obvious reasons. So, what people do is use starting prices as an instrument or as a proxy for early bidding activity, with the idea that lower starting prices lead to -- you know, presuming that starting

price is at least roughly endogenous, that leads to more bidding activity early on, and that's going to drive later bidding activity.

So, we're not -- you know, throughout all this, we're not sort of taking a stand on whether we agree or disagree with it, but we're just saying this is what's been done, and we're going to sort of work in that model. Okay.

So, basically, this test that we're going to do builds on basically what previous people have done, and the idea is if we find that low starting price auctions outperform high starting price auctions, that's evidence at least for one of the behavioral effects, because they all act in the same direction. If we find no effect, that's evidence against all of them together, okay? So, we're not identifying any one effect. If we find a positive effect, we're just saying something about all or nothing basically, okay? So, these are -- you know, without going into details, these are the papers that, generally speaking, find pretty large effects of starting price.

So, what we do is -- you know, we have some questions about sort of the exogeneity of starting price in a field setting, you know, so there's a lot of theory that says the starting price or reservation price is a

function of all kinds of stuff. So, what we do, to get a clean sort of set of data, is to run an experiment where we sold matched pairs of auctions on eBay. So, we have two identical auctions -- otherwise identical auctions, same movie, same start and end time, same sellers, same wording, same everything.

The only difference is one has a low starting price, say 99 cents, and one has a high starting price, seven or eight dollars, nine dollars, for a DVD. And basically, it's a horse race to see which one outperforms the other.

Now, obviously, we are going to wait until -- you know, we have to -- the low starting price auction has to at least be bid up to the starting price of the high starting price auction for them to be comparable, so we're not, you know, dealing with truncated distributions, but essentially, it's a pretty simple test, all right? So, we are going to use these within matched pair variation in starting prices as our test, okay?

We're also going to collect a set of observational data, and that's -- so, it's basically the same sort of data, but it's historical sort of data that we just downloaded from eBay. It's more complicated than that, but we downloaded it from eBay, and

essentially, the idea here is that it -- so, it not only lets us sort of corroborate our experimental design to make sure there is no sort of issues there, but it also lets us test this -- the Lee and Malmendier irrational limited attention idea, because we didn't have then the sort of fixed-price stuff in our experiment.

We also used it to sort of better understand some of the previous papers, which rely more heavily on the observational data, okay?

Okay. So, the first test, pretty straightforward, I think. So, actually, so there's a whole bunch of tests in the paper. I'm just going to show you sort of the simplest one for the sake of time. So, we compared the ending prices within the matched pairs of the low starting price auctions and the high starting price auctions, and we required that the auctions exceed the high starting price. So, we just -- we're not -- we're sort of comparing apples to apples.

And just for -- you know, we also run this as a more -- sort of more sophisticated, sort of left censor-dependent variable model, but it gives the same results, just in case you have any questions about these assumptions here.

So, here's sort of the first set of sort of straightforward results. Let's see how this works.

Okay, great. So, these are all the DVDs in our experiment. So, these are sort of best-sellers from Billboard Magazine, and this is -- requires both of them to exceed the high starting price, which is why this isn't the full set here. And this difference is the difference in the ending price between the low and the high starting price, okay?

So, essentially, if this is positive, then this would be evidence for the behaviors, behavioral theories, and if it's negative or null, that's inconsistent with the behavioral theories. And so basically we find not a whole lot of stuff here in the negative direction, and, you know, this is all of them together, we find the negative effect, somewhat driven by this -- there is one outlier here, but it is still negative when you get rid of it.

So, anyway -- and we can reject a positive result, okay? So, that's our sort of first (inaudible) against.

And just -- you know, in the paper we -- you know, really, to be credible, we have to explain the difference with the previous work. So, in the paper, we have a whole section here saying, you know, this is why we think we get these results, and it's different than the other results, and so there's more of a discussion

there.

So, the second set of results is about comparing auction outcomes to these fixed-price options, right? So, this is -- you know, this is Lee/Malmendier work that just came out in AR a few months ago. So, essentially -- you know, so this is a test basically from their paper, which says that -- you know, to understand their behavior, auction ending prices and expectation should not be higher than these fixed-price options, because the fixed-price options represent a reservation price on some level. But under irrational limited attention, you should see -- you know, it's possible to see auctions ending systematically above the BINs, okay?

So, anyway, so they find very large effects. So, the majority of auctions end above the BINs, really amazing stuff, and, you know, on average, auction ending prices are 10 percent above these fixed-price options, on average, right? So, really, potentially big mistakes being made by bidders in systematic ways. And keep in mind, these are sort of experienced participants, on average.

Okay. So, we had this data that we had collected for the other reason. It was an easy thing for us to check, and we were sort of aware that, you

know, through our experience, that the eBay search results are actually quite sensitive to which search terms you use, and there's also -- you know, you type in a popular item into eBay, and you get hundreds, if not thousands of items.

So, really -- you know, we had some questions about, you know, yes, maybe there's overbidding, but to assume that all bidders are sort of consciously aware of all items simultaneously and this is irrational behavior as opposed to just sort of frictions was a question for us, and, you know, it was easy for us to test. And so we thought we would check it out.

And so, anyway, let me tell you a little bit more about the algorithm then before I tell you about our test. So, eBay's search algorithm is a -- you know, they call them all words/any order algorithm. So, basically anything you put into your search box when you're searching on eBay, any of those words has to appear, proximately speaking, in the title of the listing for that listing to appear in your search results, okay?

So, if I type, you know, "Batman Begins" or "Batman Begins DVD" into my search listings, then all three of those words have to appear in the listing title for it to appear in search results. You know, you can

see the big differences you get. So, if you type "Batman Begins DVD," you know, this is what I found last summer, $6.99, and if you do -- include the word "2005," that cuts it by, you know, two-thirds or half, at least, just because the titles don't contain 2005. And then if you include the word "on," you know, you get down to five listings, right? So, anyway, the point is that it's, you know, quite sensitive to some specifics here.

You know, the flip side of that is, you know, any title that includes a lot of words in the title is going to appear in a lot of search results, obviously, right? So, if you put the word "new" in your title, you know, anybody who's typing "new" into their search results is going to find our listing, but if you don't -- let's see, how does that work? I guess the opposite. So, if you include "new" in your listing title and -- let me see. I want to make sure I don't get this backwards. Is that right? So, yeah.

So, anybody who types "new" into their search string is going to find this one, but they're not going to find this one, okay? So, basically, the more words you include in your listing title, the more search results is going to appear.

So, anyway, so this -- you know, to us, it seems pretty simple, but, you know, we -- our take-away is

that bidders -- there is going to be a lot of variation in terms of which sets of auctions and listings and BINs are going to appear in different search results, and so we suspect that frictions actually could be important.

The question becomes how you actually, you know, say this a little more formally. So, what we do is actually -- well, I'll tell you in the next slide or two.

Okay. So, this is the first set of results regarding overbidding with respect to BINs. So, these are from our observational data, and we're comparing -- you know, for each of these titles that we collected, best-sellers, we compare the fraction of auctions or the number of auctions that exceed the BIN price, right, the lowest available BIN price, and our numbers are actually -- you know, it occurs with some regularity, but it's a little bit lower than the 57 percent they find. Now, this is only for DVDs, but this is less than half the rate they find.

But maybe more importantly, when you compare the actual ending prices, there's a pretty big difference in the opposite direction; that is, auction ending prices appear to end significantly below the BIN prices, on average. So, that's -- you know, qualitatively, that's an important reversal, we think, of that result.

Now, this is only DVDs, and it's -- you know, it's -- you know, it's not a huge data set. So, this is preliminary, but then we wanted to look at this a little more carefully, okay?

So, this is looking -- you know, so we do find some -- still some overbidding here, and we want to see if we can explain this with frictions as opposed to some kind of nonstandard or irrational behavior. So, what this colorful chart is telling you is -- so, consider this here. So, this -- one indicates -- these are BIN -- auction BIN pairs where the auction contains the word "new" and the BIN does not, okay? So, these are the cases where the -- you know, the auctions are more likely to show up in search results and the corresponding BIN is less likely to show because of the wording difference. And we see that these are much more likely to be overbid. This is the overbidding right here.

Now, when you have the same words appearing, you see a lower overbidding rate, and when the BIN contains the word and the auction doesn't -- so that the BIN is more likely to show up in search results -- then you get the lowest overbidding rate. So, you can do this for, you know, all the common words you'd see for DVD, and it's -- you know, the patterns are pretty clear. And if

you combine all the words together, you see this very clear pattern.

Two minutes? Okay.

Anyway, so this tells us that not only is the rate lower, but frictions appear to be important, and so maybe sort of irrational stuff is not the whole story.

All right. So, what we did, we also took a look at Lee and Malmendier's data to see if the same kind of frictions show up there, and lo and behold, yeah. So, we find the same patterns there. So, they look at a -- you know, there's a cross-section of 12 products there. We used the word "new," which is sort of common across all products. You know, "disk" would not be relevant to, say, hair dryers, right? So, we look at new, and, again, we see the same kind of patterns for overbidding. So, when the auction contains the extra words, we see more overbidding.

So, anyway, this tells us -- this, to us, says that frictions are likely to be important. Also, you know, this isn't really the forum to point it out or to really discuss it in much depth, but there is also -- we found some sort of outliers, some data coding issues that also could help to explain some of their overbidding compared to ours.

Okay. So, given time, I won't discuss this last

one, but let me just wrap up here and say, so this -- you know, we've identified, I think, most of the nonstandard behaviors or at least many of them that have been sort of attributed to bidders and auctions, and, you know, we took a circuitous route, but we found that they tend not to -- you know, we took a closer look. We find uniformly no results.

So, this is sort of a negative paper, but, again, this is not saying that, you know, bidders are perfectly conforming to standard behavior, but we just can't reject it based on tests that have been used before.

So, anyway, so that's what I have. Thank you.

(Applause.)

DR. ASKER: So, I think about this as essentially about meta-identification in the sense that it makes me think about to what extent -- what inferred behavioral biases in the data being attributed where, in fact, what we're observing is failures in identification of just a classical model.

This seems an important question in the context of a field which I think appreciates the fact that human beings don't all the time exactly conform to the neoclassical model, but it's struggling with working out which behaviors are important and in what contexts. So,

in reading this paper, I found it very helpful insofar as it helped me formulate the standard of evidence I want to demand before I deviate from that classical model and put in some, you know, behavioral structure.

So, with that introduction, let me sort of go through the paper in a little bit more detail to show you the ways, you know, I found it provocative.

So, the research question is simple. Does bidding behavior in eBay auctions deviate systematically from the standard model? And this is a question which, as Henry said, has been asked by several people before him. His innovation is going to be to formulate a field experience that identifies very clear identification of this effect, should it exist, and it's compelling in that regard.

And so in thinking about why this is interesting in a slightly more specific context, you know, it may help us think through when behavior might invite some sort of paternalistic intervention that aspects of, say, consumer protection regulation might suggest. It provides us, as does much experimental work, a wind tunnel test of when theory is working and when it's not.

And, you know, what I think comes out particularly strongly is it provides something of a critique to aspects of the behavioral economics

literature at some sort of methodological level.

Now, in making that comment, let me be clear that it's very, very obvious in the auction context that there are instances where behavior does deviate from the standard model. So, those of you who are familiar with the Kagel, Levine and Harstad work, also some of the work by Vernon Smith, will note that in at least two instances, behavior in auctions is different from what is in the standard model. These are, as Henry said, in terms of overbidding in the second price auction, and, you know, some interesting stuff in the third price auction, and also, in the context of the winner's curse insofar as, you know, if you run a jar of coins auction, you're always going to make money off that auction, and, you know, work by people, among others, Max Bazerman, suggests that that's very hard to get people to learn not to do, suggesting some kind of problem with how we, as human beings, do certain forms of conditional probability computations in our head.

But that's not really what's at issue here. What's at issue here are the following nonstandard behaviors that have been suggested by much more recent literature, some of which is in the AER and other parts are in marketing and sort of marketing psychology journals. So, I grouped them into two groups.

The first four are part of one group really, inviting things like nonrational herding, which is the idea that more bidders must mean the quality's better, even if by construction, those more bidders can't be informative; auction fever, which is it's just all so exciting, I'm going to bid more; some sort of quasi-endowment effect, which is I've been bidding so long, I feel like I own it already, therefore, I'm going to keep bidding longer; and escalation on commitment, which is something like I put so much work into working out how to bid that I want to justify that sunk cost by bidding more.

And then there's limited attention stuff, which is a scientific way to say I appear to wear blinkers when I surf the Internet, and so I can't spot a good deal even when it's right in front of me.

So, I'm going to focus on these first four. The irrational limited attention stuff engages really with this Lee and Malmendier work. I encourage you to read the paper, if only for that part. The critique is convincing and somewhat shocking. So, let me talk about something else.

All of those four behavioral assumptions up there imply, among other things, that a lower starting cost should imply more activity when the price is higher

than a higher starting cost. Price, cost, same thing. It also suggests that the expected revenue should be higher when the starting price of the auction is lower.

So, what the authors do is run a field experiment on eBay where they used matched pairs of movies. And so what they have is an environment where the only thing that differs is the starting price of the auction. So, there will be two treatments for each movie, one with a low starting price and one with a high starting price. So, the question is, why do you need to do this?

And the reason is that it's a very convincing way to control for demand unobservables. The reason you want to do that is because the starting price is effectively a publicly observable reserve price, and that should be correlated with things that we, as the analysts, don't see but which everyone else in the environment might, all right? But through this controlled experiment, you get rid of that endogeneity problem, and the punch line is that this effect goes away, the effect that there appears to be some kind of effect of the lower starting price that raises revenues higher and bidding activity.

One of the things that I commend the authors of doing in this paper is they then take these kind of

things and then go, well, what on earth is happening in the observational data? And they go carefully through the observational data and do a much better job of identification on that data, and they find that, you know, they can show that when you do the job properly, the effect goes away.

So, let me get to my comments, which are -- I mean, all of them point in the same direction, which is this is a courageous paper to write, because in contrast to much of the behavioral stuff, it's really saying, at least in this set of environments, there's nothing there. So, if I was advising a graduate student on how to get published in the AER, certainly in this field, this would not be the way I would advise them to do it, all right? I'd advise them to find something that was startling and thought-provoking and at least positive and write a paper about it.

The problem with that is that as we start to think about the extent to which we want to adjust our models to take into account different forms of behavioral bias, having literature that's only full of positive results could well be misleading to us. And my concern with the evolution of the field is that whether by accident or possibly by design, there's some kind of publication bias that would lead us to infer a greater

prevalence of these kind of behaviors than actually exists.

So, you know, in saying that this paper is courageous, I think that's probably right, but I think it's also important in the context that it provides a welcome counterbalance to a lot of this literature that's purporting to find effects, both just in showing that these effects aren't always there and also in providing us some basis on which to think about the kind of evidence that we need to accept the presence of these things in any given context.

Thanks very much.

(Applause.)

DR. ROSE: So, I am sensitive to the fact that we're running 11 minutes late, and I don't want you to be antsy to finish before you've heard the third exciting paper in this trio. So, I think given that John has given you such a great window into the previous paper that I hope maybe we'll bring Henry up -- and I don't think we'll have time for questions for Mara's paper either -- so we will bring the authors who don't have time for questions up to the front, and you can congregate or congregate as we move out to the reception area. I am going to move right to the third paper.

So, Mara Lederman.

DR. LEDERMAN: I just need to find my slides.

DR. ROSE: I can help you.

DR. LEDERMAN: And I am still going to take 20 minutes?

DR. ROSE: Yes.

DR. LEDERMAN: All right. But I won't take more than that.

DR. ROSE: And this will be discussed by Jeff Prince, and we'll have Jeff just go right after.

DR. LEDERMAN: Okay, great. Let me just highlight two things before jumping in. Number one, the title has changed since the title was originally circulated under, which is the one that appears on the program, so I hope there's no confusion there. The work is joint with Silke Forbes, who is here, and Trevor Tombe, who unfortunately couldn't be here today.

I know everybody's tired, it's the last paper, it's quarter to 5:00. The good thing about this paper is, one, it talks about an industry everybody knows about, which is airlines; two, it talks about a feature of this industry everybody likes to complain about, which is airline delays; and three, everything I'm going to convince you of, hopefully, comes out in very simple pictures of raw data, basically. So, you won't have to think that hard, I don't think, at the end of the day to

buy what we're going to try to tell you.

So, the paper is about quality disclosure programs and gaming, and specifically we're going to be thinking about the incentives that employees have to carry out gaming. So, let me start by motivating the paper.

So, as everybody knows, as we've sort of hinted at in some of the talks already today, quality disclosure programs are a big deal. We're seeing them in more places, and the objective of these programs is to provide systematic information to consumers about product quality in settings where we don't think consumers are well informed. So, we see these in the healthcare setting, let's say hospital report cards; we see this in education, with student test scores; we see this in the restaurant industry with hygiene scores.

There's been, you know, a growing amount of empirical analysis of these programs. In general, they find that they work in the sense that firms seem to be improving quality in response to the introduction of some form of disclosure, but there's also a growing amount of evidence that firms game the programs in some sense, and we're going to sort of loosely use the term "gaming," and I think the literature has as well, to refer to, you know, an effort to improve quality on

dimensions that are reported, potentially at the expense of quality on dimensions that are unreported.

And if the reported measures are only imperfectly correlated with what consumers actually care about, then this kind of gaming behavior may both lead firms to allocate resources to the wrong places and distort the information that consumers are seeing. And I think -- you know, I think it's possible, and after sort of hearing everything we've heard today, especially in the first session, to believe that, you know, when we disclose information, we're not disclosing everything consumers could care about, either because consumers are heterogenous, and so what we're disclosing was what some people care about but not others, or perhaps, more importantly, because when you disclose information, you face a trade-off between disclosing lots and lots of information that probably no one's going to pay attention to versus disclosing just very simple information that may not capture everything.

What we want to focus on sort of relative to the earlier literature is thinking about or anticipating when is gaming going to happen and really thinking that gaming won't just depend on the design of the program. That's one thing that's going to be important, but it's also going to depend on the characteristics of the

product and the incentives in place in the firm.

And in particular, what we are interested in -- I just realized I should be talking into this -- you know, the questions we're going to ask or we're going to think you need to think about is, well, what are the dimensions of quality that a program is trying to measure? How can those be manipulated? Who's in a position to manipulate them? And do those people who are in a position to manipulate them actually have incentives to do so? And that's really what we're interested in thinking about in this paper.

So, let me tell you what we do. We are going to investigate the relationship between gaming of a disclosure program and the incentives provided to those employees who we think are most likely to have to carry out the gaming. So, basically, in a sense, we have sort of a disclosure environment that's held fixed, you know, for a long period of time, but we're going to have cross- and within-firm variation in the explicit incentives given to employees based on the firm's performance in the disclosure program.

Our context is this airline -- you know, is airline on-time performance. I'll tell you more about that in just a minute, but we think these issues are relevant in other settings as well. What you need to

know really quickly is that the Department of Transportation, who collects and disseminates information on airline delays, counts flights as being late if they arrive 15 or more minutes after their scheduled arrival time. That's sort of the disclosure program. And based on this, the DOT is going to create monthly rankings of airlines, which will get picked up in the media and you have probably seen before.

Four useful features of this setting that I think lend itself to exploring the kinds of things that we're really interested in: Number one, from a design perspective, we think this program makes it very clear what you need to do to sort of game the program. Don't have flights that are 15 minutes late, right? I mean, those are kind of the worst flights to have. If they are just one minute earlier, they count in the on-time column, and it probably wasn't that hard to shave that one minute off.

But what's interesting here is that airlines can't predict in advance which are going to be those 15-minute late flights, right? Maybe you know which are going to be really late, because they fly to congested airports, which aren't, but you don't know what's going to be a 13-minute versus a 15-minute versus 17-minute late flight.

So, if you are going to game in the sense that you reduce those 15-minute delays, it has to happen in real time, when you realize this flight is sort of the candidate for gaming.

What's very cool and interesting or puzzling -- you know, pick your favorite adjective -- five airlines of the big airlines, over time, have implemented firmwide employee bonus programs based on the airline's rank in this government program. They all face free-rider problems, because they're firmwide, they cover all the employees, but they differ in the incentives they provide, because they make it harder or easier to achieve the target on which the bonus is awarded.

And as you know, if you've seen airline papers, we have access to lots of data. In particular, we have millions and millions of flights that we can look at, and that lends it -- you know, lets us set up what we think is a pretty clean identification strategy, which I'll tell you about in a minute.

So, let me preview the findings in case we do run out of time or in case you have no more energy for listening after this slide. Number one -- and I am going to thank John for sort of, you know, touting all the -- you know, the importance of believing null

effects. Number one, we find no evidence of gaming by airlines who don't have these bonus programs in place. So, despite the fact that the program creates this clear incentive, we just see nothing different about their behavior with respect to these 15-minute flights.

We see no evidence of gaming by sort of three of the five airlines who introduced programs when they introduced programs with targets that couldn't realistically be achieved, and I'll show you what we mean by that.

Two of the airlines who do introduce programs with sort of realistic targets, for them, we see very strong and we think convincing evidence of gaming, and I'm going to show you that.

Here's an example. It comes right out in the raw data. So, this is just a histogram of Continental Airlines' arrival delays before they introduced a bonus program. The red line is 15 minutes. You can see sort of, you know, a lot of mass around zero. Then they introduce a bonus program, and they have a lot more flights that seem to arrive exactly 14 minutes late and fewer that arrive 15 and 16 minutes late.

So, kind of what you need to know is two airlines' data looks like this and the other three don't. And now I'm going to just sort of convince you

of it in a slightly more rigorous way.

So, let me tell you a little bit of background on the disclosure. This program goes back to 1987. The rule is basically if an airline accounts for more than 1 percent of domestic passengers, they have to report their scheduled and actual arrival and departure times. Over time, as big airlines have shrank and small airlines have grown, more and more airlines have qualified for this reporting requirement, such that the set of airlines you're ranked against has grown from basically 10 to 20, peaked at 20, now 16.

A flight is late, like I said, if it's 15 or more minutes behind schedule. We get these rankings based on this. These go into the media, and we have evidence out there that consumers do respond to sort of -- you know, to delays. It's something they care about.

Something that's relevant -- we may not have a lot of time to talk about it -- but people want to know, how does this data get reported by the airlines? Can they just lie about when a flight got there? The answer is yes, and sometimes -- yes, they can, and sometimes we think they do.

So, basically, there's two ways, historically, this data could be reported. Either it was reported

manually, where literally someone at the airline recorded the actual arrival time and it was compared to the scheduled one; or increasingly and now entirely, it's reported automatically because the airplane has a given technology that just sends information, like, you know, the engine is off, the door is open, whatever it is that conveys, you know, the flight has arrived at the gate.

During our sample period, some of the airlines, especially those that have bonus programs, are combination reporters in the sense that some of their planes have the technology and some don't. We don't know which planes are which and which, but we have an approach to try to pick up the manual planes, and we think it works pretty well, because we see bigger effects on those planes. So, we think some of what's going on in lying. We don't think that's everything.

So, here's a little more -- here's a little more raw data, just to give you a sense that the manner in which they report seems to be important. For airlines, once we know that they're fully automatic, you see their histograms are very smooth. For the airlines who are fully manual -- and that histogram is sort of dominated by Southwest, it's by far the biggest airline in there -- it's not smooth. There seems to be a lot of

rounding with a lot of delays -- a lot of flights being zero or five or 10 minutes or 14 minutes late.

The combination reporters -- so, these are guys who in this time period we know have at least some manual planes -- they look pretty automatic just in terms of what their histogram looks like, but we see they seem to round at zero and do something at 14 minutes.

So, let me just tell you a little bit about these bonus programs. Basically -- I won't have time to go through the details, but they pay a fixed amount, between 65 and 100 dollars, to every employee in a month in which the airline meets whatever their target in the ranking is. What's interesting to take from this table is the heterogeneity across the programs, and so there's two things, I think, that are worth noting.

The two early programs -- and these are the ones where we are going to see evidence of gaming -- rewarded the top five spots at a time that only ten airlines were large enough to qualify for the rankings. The three later programs rewarded only first or sometimes first and second at a time that between 17 and 20 airlines qualified for the rankings, and among that set were some small airlines, for example, Hawaiian Airlines, who always ranked number one in every single month after it

qualified, because it flies in good weather to uncongested airports.

And so in some sense, you know, I read that if I work for U.S. Airways, the chance, when my average rank is about 9.8 in the year before, that I'm going to, you know, get my $75 if we're first against Hawaiian is basically zero, right? So, we think -- you know, we can't test it formally, but that's what we think is explaining why we don't see any action in response to those programs.

So, let me talk you through the empirical approach, which is actually -- it's pretty intuitive. So, what do we want to measure? We want to basically measure, do airlines systematically try to reduce delays on flights that they think are going to arrive right around 15 or 16 minutes late? So, we need to do three things.

We need to find those flights that look like they're going to be 15 minutes late. We then need to look at sort of what happens on those flights. And then we need to say what would have happened on those flights so we can say that what happened is, in fact, what we call gaming; is something that we wouldn't have happened absent the incentive to get them under 15 minutes. And I think our setting lets us do all three in a pretty

straightforward way.

The first thing we can do is for every flight, we observe its progression through the stages of its flight. So, when it left the departure airport -- when it left the departure gate, then it taxis, it leaves the departure airport, then it flies, then it lands, then it taxis in. So, what we can do is we can take a given stage of the flight, and the stage we focus on is when it arrives at the airport, we can look at how delayed it's been.

We know all the delays that have been incurred so far. We can estimate sort of what would typically happen afterwards, and we can calculate a predicted delay for a flight, where everything except the final stage, the taxi in, is based on sort of data that's happened already. So, that will allow us to pull out the flights that we predict and the airline probably would predict are around 15 or 16 minutes late.

Then what we can do is just estimate whether subsequent delays after that, which are only going to happen through taxi-in times, are systematically different for those flights that we think are near the thresholds. And then in terms of looking at a counterfactual, in terms of looking at --

Is that zero to me? Oh, okay. Sorry, that

threw me off. She's holding up a -- okay.

And then in terms of looking at a counterfactual, we can exploit the discontinuity and say, "Look, these 15-minute flights shouldn't really look any different than 13- or 18-minute flights, and they certainly shouldn't look different than -- shouldn't look better, in a sense, than flights that are really, really late, that say, 20 -- predicted to be 25 minutes late." So, that's what we're going to do.

And so I'll just give you an example of how we do it. Let's take a flight, Flight 236, you know, by Delta, between Boston and Atlanta, some month. Suppose it was supposed to get to Atlanta at 4:30. If its wheels -- and that's at the gate. The arrival time is at the gate. If its wheels touch down at the runway at 4:36 and the median taxi time, the typical taxi time this flight has in a quarter is four minutes, then we'll predict that it will get to the gate at 4:40. It was supposed to get to the gate at 4:30, and we predict that it's ten minutes delayed.

So, we are going to do that for every single flight in our data, and we are going -- we have done it with sort of exploring the robustness to different ways of predicting the taxi-in time, and then we're just going to make dummy variables that capture how predicted

this flight is to be delayed, ten -- how predicted -- you know what I mean, 10 to 11, 11 to 12, 15 to 16. We'll put less than 10 in a bin, more than 25 in a bin, and we'll make these dummies for every -- for all the airlines who don't have programs, we'll put them together, and for every airline that introduces a program, we'll look before and after, when we can, so we can just estimate sort of a whole bunch of coefficients on these dummy variables and see who's doing what.

And so then we're going to estimate regressions where we regress a flight's taxi time on its predicted delay. We'll put in a bunch of control variables, and we're going to put in fixed effects for the carrier arrival airport day. So, we're going to take all of, let's say, American's flights that land at Reagan today. We're going to calculate their predicted delays. The variation there is driven by stuff that happened before they got to Reagan, right, because these are delays incurred before they land. And then we're just going to look if those we predicted to be 15 minutes late sort of have shorter taxi-in times. And we'll do some stuff with, you know, standard errors and all kinds of things that I won't talk about right now.

So, let me show you sort of -- this is -- our first set of results -- all the results I'm going to

show you right now are all out of one big regression, but I'm showing you sort of groups of coefficients together.

So, these are the airlines in the early time period, '95 to '98, the time of the Continental and TWA programs. These are the airlines who have no bonus programs. Lots of numbers. What you should look at is see that these numbers all look the same. And so basically what a coefficient tells you is approximately the percentage change in taxi-in time for a flight with that level of predicted delay compared to a flight that's predicted to be less than ten minutes late. And we see all these flights are sort of 3 to 4 -- about 3 to 4 percent shorter taxi-in times.

When we look at Continental, we can't look at them before their bonus, because we don't have the data on taxi times before '95, but you what their histogram looked like just on sort of raw data before, and we see systematically shorter taxi times for precisely those flights that are predicted to be 15 to 16 or 16 to 17 minutes late. So, their taxi-in times are about 14 percent shorter or close to a minute shorter, somewhere between 45 and 60 seconds shorter.

When we look at TWA, for them, we can look both before and after their bonus program. We don't see this

pattern for them before their bonus program, but then they show a very similar pattern to Continental and actually pretty similar magnitudes after they introduce their bonus program. So, the easiest way to see it is to just plop the regression coefficients, and this is what we mean by nonmonotonicity.

You might expect the taxi times get shorter for flights that we expect to be later, but we don't expect them to then get longer again for flights that are 17, 18, 19 minutes late.

So, that's what we see in the data, in the raw data, and again in the regressions. When we run analogous regressions, doing predicted delay in the same way, looking at the three later programs -- these come from different samples. We can't put all the data from '95 to 2010 in a single regression. There's just too much data. As it is, we have to sort of randomly sample our flights. We don't see any evidence of this kind of behavior for any of these three.

You know, for United, we don't see any sort of differences in taxi-in times. For American, you know, relative to the sort of on-time flights, taxi-in times are 4 to 5 percent shorter. U.S. Airways, again, we just don't see the nonmonotonic pattern. And in all our subsequent empirical analyses, we don't see any evidence

of this happening for them. So, I'm not going to sort of revisit them in the later analyses, but all of our analyses we've carried out on all of the programs, and we never see effects for these three.

Okay. So, the question we could ask, and this is just sort of a another way to look at the data, saying, when they do this, does it work, right? So, if they try to speed up these flights, do they actually get it there under the threshold? And to look at that, we are going to take the exact same regression but just replace the left-hand side variable with a dummy that equals one if the flight gets there exactly one minute earlier than we predicted. We'll do that for every flight in the data, and in a sense, what these regressions do is they test whether we are systematically worse at predicting delays for flights in that critical range, right?

So, we have all these flights. We predicted when we thought you'd get there. For some reason, are we worse at our prediction for those flights? Are they more likely than any other level of flight to get there one minute earlier than we predicted? And we do the same thing for getting there two minutes earlier.

And so the coefficients in these regressions tell you basically the change in the flight's

probability of being one or two minutes earlier than they were predicted relative to flights in this less than ten minute bin. So, I'll just show you plots of these coefficients.

What do you see for Continental? You know, by far, flights that we predict to be 15 minutes late are -- you know, are much more likely to get there exactly a minute sooner, right? Ten percentage points more likely. On average, flights in the data get there one minute earlier than we predict only 20 percent of the time. So, this is a really big effect.

When you look at the two minute earlier, it's the same thing. It's these 60-minute flights, you know, or whatever, there's 11 or 12, might be 13 percentage points more likely when the average flight for Continental only gets to the gate two minutes sooner than predicted 10 percent of the time. So, these are really, really big effects. TWA's pictures look similar, but for time, I didn't show it.

So, then, a little bit about these manual planes, because it's kind of interesting. It's important, we think, regardless of whether they're lying or speeding up planes, it's interesting that they're responding to this incentive, but we would like to know something about how they're responding. And so

basically, let me tell you intuitively what we try to do.

If you remember those histograms I showed you at the beginning, what you see is that carriers who report their data manually seem to have a lot more flights that arrive with exactly zero delays. And so what we're going to do is we're going to take a plane, we can track the physical plane, and we're just going to calculate basically how often in a year it arrives with exactly zero delays, and if it arrives with zero delays sort of too often to be sort of what someone who reported their data automatically would find, we're going to flag it as a manual plane. Does that make sense?

And then we're going to separate those planes, create these bins separately, and see if things look different. And so we're going to do that for Continental and TWA. And when we do that, you can see -- I'll just show you for Continental -- again, TWA looks similar -- the effects are larger and more pronounced for the manual planes. Those are the ones that are in red here, but the same pattern exists for the automatic planes. We still sort of pass our hypothesis test.

We still see that behavior for the 15-minute flights is different than the one or the 18 or the 25,

and the magnitudes of the effects on the automatic planes are pretty similar to what we estimate in the full data, just because there's not a lot of manual planes. It's about 10 to 20 percent at most of their fleet.

And this is taking a pretty conservative approach to basically calling anything that we think could even look manual as calling it manual, so we can try to have as clean and automatic sample as possible. So, we interpret this as saying some of what we're measuring is probably lying, but some of it seems to be actually shorter taxi times.

So, since I've got about a minute left and maybe not even that, let me just sort of highlight, since we can't -- this is something we can't get at sort of rigorously in the data, is why do we think we see this response for the early programs and not the late? And I think we have enough data and a good enough identification strategy to believe that sort of the null effect on the late programs really is sort of no effect.

I don't think we're missing it. I think it's two possible things. One is this misreporting. So, in the later time periods, those guys are reporting automatically, there is no scope for lying, and lying seems to be part of the story here. But more

importantly, we think the later programs just have much weaker incentives.

The chance of hitting first or second when you're ranked out of 20 and that 20 is just a more heterogenous set of airlines, including regional carriers and low-cost carriers, just provide very, very sort of weak incentives for the employees at that those airlines.

So, sort of summary and concluding thoughts. We think -- when we started this paper, we thought for sure we'd find tons of evidence of this gaming everywhere. We thought the program design was so obvious to encourage this manipulation around 15 minutes. We think it's interesting that we don't find evidence except for when these explicit incentives are provided to employees.

On the other hand, it makes sense in this setting, right, because this can't be a top-down thing. Someone on the ground has to know that this is sort of a critical flight and exert the effort, and that someone probably needs to have the incentives to do it, right? I mean, on-the-ground airline employees don't have any indirect way of caring about the airline's rank in this government program.

We have done some simulations and then we show

you that just by reducing the selective reductions in taxi-in times can buy you one to two spots in your rank. It doesn't move you up five spots, but it can move you one to two spots.

So, just our concluding thoughts on sort of what we see as the contributions of this paper. We think the paper contributes to this growing empirical literature on gaming of disclosure programs. We think it starts to bring that sort of information, economics literature and org econ literature together a little bit. We think it's the first to explicitly consider the link between gaming and the incentives provided inside the firms who may be trying to game a disclosure program.

It highlights the importance of not just thinking about program design but also sort of how is quality produced and who's producing that quality and what their incentives are. You know, as is obvious to everyone here and has become sort of even more obvious during today's talks, this is -- you know, how we provide consumers with information, how they use it, what they can use and what firms will do as we change disclosure is obviously part of a policy-relevant debate in a bunch of important settings.

And so hopefully we see these -- you know, our work as contributing to that, and we think it's just

kind of neat that at least we find that these really high-powered incentives don't seem to change behavior precisely because they're just too hard to achieve their reward.

With that, I will end hopefully on time.

(Applause.)

DR. PRINCE: Okay. I'm in the enviable position of standing between you and hors d'oeuvres, so I'll go as slowly as I can.

So, thank you for having me give this talk. For many reasons, I think Mara already communicated how interesting this paper really is. It was a pleasure to read it. As someone that's dabbled in some structural stuff and reduced form stuff over my career, I can say it's really nice when you get this rare moment that you can just see what's going on in the before and after picture, right, and it's a real after picture. It doesn't have a sun tan. You know, it didn't have slimming clothes on, right? This is the real deal. And Mara showed you that, right? So, I could see that right when I was reading the paper. It was really clean, really nice to look at.

So, as I mentioned, there's a lot of other reasons to like this paper. I'll dive right into some things that I think might be worth adding as far as my

comments are concerned. I put them in three broad categories: General gaming insights, airline insights vis-à-vis on-time performance, and possibly just throwing out a different measure of interest that perhaps could be done to round out some of their results.

So, to open it up, I -- as I was reading this, one way to kind of read through this is to think of this as being like a verification of theory, right, where you could say -- and this is -- you know, it's entertaining the theorists in the room, the empiricists trying to do some theory up here. You know, in the back of my mind, I'm thinking this might benefit from a little bit of a toy model to get things rolling, to fix ideas, because everything kind of maps into this, where you say if you think about these front-line employees, what are they doing? They say, I've got utility, it's increasing in my pay, it's decreasing in my effort, and what's happening? I'm -- my world is changing to where now my pay depends on my rank, and my rank perhaps depends on my effort, and my effort could be measured in ways of real effort, I'm hustling, or I'm cheating, right? But either way, you'd say that things have changed for me as far as how I'm going to make decisions to optimize my utility.

And the way it's structured, right, the pay-off from your effort is greatest if you know you're somewhere near the threshold, right? So, that's the way this program is designed. So, that's one way to perhaps go with this. I'm not sure if this was the authors' intended way to frame the issue.

On the flip side, you could say, you know, you've got -- you've got these -- this -- you could put the theory model in there. Another way to go is they -- they put in the introduction this kind of pseudo-motivation that you might want to be reducing delays most on very delayed flights, ones that are, you know, getting in the two-hour range. I would say that's moving up the ladder, though, if you start thinking about it that way. So, on the ground, you're probably not thinking about that very much, but I think that opens the question as to why is this the incentive scheme, right? So, why did they decide to go with these rankings that clearly motivate you to game the system?

And so I just -- I think -- you know, that is obviously not the focus of the paper, but I think it's something that might warrant some mention, because I found that a bit baffling. And you could bring out a bunch of reasons why that might be, right? So, maybe it's costly to try these other measures. Maybe they

only care about the rankings because they think that's all that matters. Maybe they believe the 15-minute margin is the right one. This is what's best for their consumers or for their profits. And maybe they're not even aware of the gaming behavior, right? Maybe they think that putting this in there, the on-the-ground employees would never want to game the system, right? No way.

So, given a large -- for the golfers in the room, you could say there's a large proportion of the reduction in delays is due to the whole pencil wedge, right, the best club in my bag? One question that it brings out there is why wasn't this happening before the incentive scheme for Continental and TWA?

So, the before and after I think was a really compelling show, but I also sat there and I thought, well, cheating seems costless, right? Why wouldn't they just have been cheating anyway before there was an incentive scheme, right? I could have just been pencil-wedging this all along. And in a sense, that tells me something. So, you could say, well, maybe -- are there consequences to tinkering when it's manual? You know, is management at risk if they come in and they say, you know, if this thing's two minutes past, just put it down as 14, it's all good? You know, that made

me think about what's going on here, and I think -- I'm not -- I don't know. You know, I've actually studied this industry. I'm not sure what the consequences are, but it suggests that there might be, the fact that they weren't doing it anymore or they weren't doing it before the incentive schemes were put in place. So, at any rate, the incentive schemes show us something about employees' thresholds for dishonesty.

So, then, if you go to insights for on-time performance, the incentive scheme obviously implies a huge free rider problem, right? The thing, I think, all of us are asking ourselves is is it plausible that the person on the ground is saying, this flight is close, right, we're close to the 15 minutes, I'm going to run down the tarmac and get this sucker to the gate as fast as I possibly can, when my expected pay-off from that one improved flight could be really, really small, right?

And so if that's not what's going on, then we say, okay, well, then, it's probably lying, right? So, it's something that's relatively costless. And I think Mara pointed out, you know, they're trying to sort out between the two, because either way, it's gaming the system, but it's important as far as what we think might be welfare effects or other types of measures.

I really love their idea about breaking it down, manual versus automatic. I just wasn't sure if that could fully distinguish the two, right? So, they had this way of trying to say these are the manual flights. If you claim it's all lying, though, it's hard to refute that claim based on that measure. So, I would just say that one thing you might want to do is if it's possible -- and it's probably not, if it was, I would imagine they would have tried it -- but if there was a way to just say these are for sure automatic planes, right, so be able to pin down this is a real effect that was happening, that would be useful not just for this paper, but I'm really interested in that, too, having done on-time performance stuff, because a lot of times people are skeptical as to whether there are real changes in on-time performance, real quality changes that happen based on competition, based on incentive programs, right? Do airlines really have that much control over their on-time performance? A lot of people just come out and say it's all congestion, it's all airport effects, these guys are pretty much just rolling with that.

So, even if it's just lying, the effects of the incentives programs, as I mentioned, they imply a cost to lying. They give us a sense of a sufficient pay-off

to induce employees to lie, because they weren't doing it before on some of these airlines, and if we consider the free rider issues, it appears to be a small expected pay-off is necessary to get them from going from not lying to lying.

However, given the employees weren't lying before the incentive change, it appears that you need something, right? I won't do it unless you give me at least a little bit of something to get me going.

The last thing -- I know I'm almost out of time -- the last thing I wanted to mention is while we're at it, if we're thinking about real changes in on-time performance, I'd be curious if you'd think about the pilots' ability to effect this as well, all right? So, you could say, well, the on the ground people know the most, right? When the wheels hit the ground, we know we're close, we could perhaps really get this thing under the threshold, but you could make the case that the pilots have the most power to make a change, right?

So, we have all been on flights that are late and the pilot basically says, well, I'm going to floor this sucker and we are going to get there way faster than our scheduled flight time, and I wonder if the pilot could sit there and say, well, we have this incentive program here, I know that we're going to be

roughly close to the 15-minute threshold, this is the time when I'm going to make an effort to pick up the plane, right, get it going.

So, that might be something that could at least be looked at. You could say, well, let's look at wheels up, wheels down as a function of predicted delay. Obviously, you're going to have less precision as far as what the expected delay was at that point, right, when the wheels go off the ground, but you could get a rough idea, right? You could say, well, let's look at the ones that are predicted to be about 10 to 20 minutes versus the ones 60 to 70 minutes. Do they seem to put a little extra hustle on that 10 to 20 because they know that that could be the one that makes the difference, right? And so that -- that would be pretty interesting, in addition to what I think is already interesting in this paper, because it would imply real on-time performance changes, right?

This is real time being saved. The welfare implications wouldn't be clear, though, right? So, what are the costs of flying faster, the fuel costs, maybe safety concerns, things like that. But I think that would be an interesting extra measure that would be relatively easy to throw in there and could get to some of the real effects that I know you're trying to tease

out, in addition to some of the pencil wedge stuff.

So, overall, I really enjoyed it. Thank you for the opportunity. And hors d'oeuvres.

DR. ROSE: So, I want to thank you. You've been a most attentive audience, and I think the authors -- if the authors just want to come up to the front, maybe you could start by making it easier for people who have questions to find you, once Chris releases us to the cocktail hour.

DR. ADAMS: Yes. So, we're going to have some food and drinks just back there, and feel free to stay, talk to the authors or anybody else. And then we'll go through until about 7:30. One issue is if you leave, you can't get back in the door after 7:00. So, note that. And our parking lot, I think, closes at 7:00. So, that's another issue. But otherwise, thank you all very much for today. It was a great session.

(Applause.)

(Whereupon, at 5:19 p.m., the conference as adjourned.)

C E R T I F I C A T I O N O F R E P O R T E R

DOCKET/FILE NUMBER: P085800

CASE TITLE: MICROECONOMICS CONFERENCE

DATE: NOVEMBER 3, 2011

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the notes taken by me at the hearing on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: 11/17/11

ROBIN BOGGESS, RMR

C E R T I F I C A T I O N O F P R O O F R E A D E R

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

SUSANNE BERGLING