FEDERAL TRADE COMMISSION

I N D E X

FEDERAL TRADE COMMISSION

FEDERAL TRADE COMMISSION	)	
WORKSHOP ON:	)	Matter No. P085800
MICROECONOMICS CONFERENCE	)	
-------------------------------	)	

THURSDAY, NOVEMBER 18, 2010

Conference Center
Federal Trade Commission
601 New Jersey Avenue, N.W.
Washington, D.C. 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. ROTHSTEIN: There are some preliminary announcements before we get going. I'm Paul Rothstein. I'm one of the co-organizers here. If you could, as always, turn off your cell phones, I would appreciate that.

The conference is being recorded, really typed up as you speak by a stenographer, so when we come to the questions, we have to insist that you use a microphone so that it's clear to her. She can hear what people are saying, but also I would encourage the speakers, if you have the choice between speaking a little faster and a little slower, speak just a little slower.

The rest rooms: You can get to the restrooms without going through security if you're careful, so you just have to go out the door you came in, and it's a little jog to the left and then you'll find the rest rooms there on your left. You will see the pathway through. You'll be passing to the left of the security desk, but you won't be going through the monitoring gates.

We have Internet accessibility here. The password instructions are at the registration desk if you need to get into your Email.

Now the security briefing: If you go outside the building without an FTC badge, you will have to get back through by going through the whole x-ray machine process again, so just a note in terms of how much time that might take you. And in the event of a fire, common sense things, you will get out of the building.

Ideally you'll head to Georgetown, across the street because you'll exit the doors, it's the only way you can get out, and you will go across that street, and you will find yourself at the Georgetown University campus. If someone is as angry at Georgetown as they are at the federal government and there's trouble there, then just proceed to some place safe.

On the other hand, there might be some circumstances where it be safer to stay in the building, and in that case you will be told where to go inside the building. It will probably be some portion of the parking garage.

There is the unusual request that if you spot suspicious activity, please alert someone.

It's my pleasure now to introduce the Deputy Director of our Bureau, Pauline Ippolito. Pauline earned her doctorate at Northwest University in mathematics. She has worked in a variety of positions at the Federal Trade Commission and in the Bureau of Economics. Her research and policy interests include the economics of risk and information in consumer good markets and the design of public policy for advertising and labeling.

She's done a lot of work on advertising and information related to health food claims in food products, and right now she and I are working on a project involving food marketing to children.

She's been involved with the Agency's fraud and ID theft surveys and in general efforts to improve consumer disclosures in many areas, but including the mortgage markets, which is a very hot topic right now.

So it's a pleasure to introduce Pauline Ippolito.

(Applause.)

DR. IPPOLITO: Well, first welcome. We are very happy to see such an interesting crowd, and I think you will agree, the program looks quite promising today, so we're all looking forward to that.

Our Director at the Bureau would have been the one introducing you and welcoming you to the conference. Unfortunately, he has to go and defend the merger guideline revisions in front of the ABA, so that's where he is this morning, but he'll be joining us later.

I should mention that Northwestern is our cosponsor, the Searle Center from the law school and the Center for Industrial Organization at Northwestern.

So are we going to do a program adjustment?

DR. ROTHSTEIN: We will do a program adjustment, yes. We'll do the first panel.

DR. IPPOLITO: So our key note speaker is not here, so we're going to change the order in which we do things, so let me turn things over to Tim Daniel, who is orchestrating our first panel.

PANEL SESSION ONE: Disclosures and Informed Consumer Choice

TIM DANIEL, Federal Trade Commission

PHILLIP LESLIE, STANFORD UNIVERSITY

JEFFREY BLUMBERG, TUFTS UNIVERSITY

DR. DANIEL: Good morning. My name is Tim Daniel. I'm an economist here at the Federal Trade Commission in the Division of Consumer Protection. It's my pleasure to moderate this panel, which is entitled "Disclosures and Informed Consumer Choice."

I am very pleased to have this opportunity and very pleased to be able to introduce these two gentleman to my left, both of whom have worked in interesting ways in this area in recent past.

To my immediate left is Dr. Jeffrey Blumberg from Tufts University. He has a Ph.D. in pharmacology from Vanderbilt University and is currently a professor in the Friedman School of Nutrition Science and Policy at Tufts.

He also holds the title of senior scientist and Director of the Antioxidants Research Laboratory at the Jean Mayer USDA Human Nutrition Research Center on Aging at Tufts University. I tried to memorize that, it just didn't work out.

Of particular note for today's session, Dr.

Blumberg serves as a scientific advisor to the Guiding Stars Program. The Guiding Stars Program is an in-store information guidance system in supermarkets designed to provide consumers with point of sale information on the nutritional value of their food choices, obviously designed to try to improve those choices.

To Jeffrey's left is Dr. Phillip Leslie from Stanford University. He has a Ph.D. in economics from Yale University and is currently associate professor of economics and strategic management in the Stanford Graduate School of Business.

Dr. Leslie's research agenda includes an examination of the role of information and in the behavior of firms and consumers, and he recently coauthored with two of his Stanford colleagues, Bryan Bollinger and Alan Sorenson, a very interesting paper, an empirical assessment of the impacts from a law in New York City that compels a disclosure of the calorie content of the foods offered at fast food restaurants.

They somehow got some very interesting data from Starbucks and did some very interesting empirical research on the impacts of that calorie posting law.

As to how I would like this hour to work, I will turn the mike over to Jeffrey for about ten minutes and then to Phillip for about ten minutes. Then I will speak for about the same amount of time on some topics.

Then I would like it to turn it into a round table discussion. I'll pose some questions to our panel members, but at that time, I would like to invite everybody in the room to feel a part of that discussion, and we'll recognize people to make comments, ask questions and so on.

So if I could, I would like to turn it over to Jeffrey.

DR. BLUMBERG: Good morning. So the topic is disclosures and informed consumer choice, and one approach to that is putting icons or information, numbers on the Front-of-Package. I'm going to talk briefly about the experience I've had with the Guiding Stars Program, which is now actually a licensing company selling this nutrition navigation or nutrient profiling system to supermarkets across the country.

You heard I'm a professor at Tufts University, so I'm really speaking as an academic, not for the company, but I think it's a good program, and so I don't mind describing it for you, but I do disclose that I am on their Scientific Advisory Panel and helped to develop this program, so I might be a little biased. I'm also compensated for that effort.

Guiding Stars and another program called NuVal are not technically Front-of-Package. They're on the shelf tags where you see the prices. As you know, the Front-of-Package belongs to the manufacturer, but the shelf space and the shelf tag belong to the supermarket. The intent of these two different programs are the same.

They're really quite different in where they come from, manufacturer versus retailer, and nonetheless, they are considered currently by FDA as being the same thing. They're Front-of-Package labeling.

When one labels the shelf tag, as you can see here and in Guiding Stars, it's this icon of a little running man and either one, two or three stars or no icon at all. Consumers seem to understand right away what the intent was. One star is good. Two stars is better, and three stars is the best in terms of making a nutritious choice.

These tags also show up in the produce department on prices or elsewhere on the shelf and for things like deli products and meats and other produce where the labeling, like on this oven roasted meat item, will be on the weight UPC tag.

So Guiding Stars is a balanced approach to nutrient assignments. It is based on consensus science; that is, guidelines that have been established by either Dietary Guidelines for Americans from the Department of Health and Human Services and USDA or by FDA regulations and guidelines. Where we couldn't find adequate help from those, we went to the World Health Organization, American Heart Association, and so on to create the algorithm.

Very simply it is a balance score of positive nutrient attributes: Vitamins, minerals, dietary fiber and whole grains and nutrients to limit the negative attributes like trans fat, saturated fat, cholesterol, added sugars and added sodium. We get this information from the nutrition facts panel on the food products except where there are none, for example, on fresh produce, and then we use USDA databases to develop them.

We decided, and I would tell you other programs are quite similar, on a nutrient density approach; that is, it's based not on serving site, not on package weight, but on one hundred KCALs, calories, so it doesn't matter what the size of the package is. It doesn't matter what the recommended serving size is. They're all based equally, and this takes care of both weight and volume issues, and we just think it's a reasonable way to go.

It's a tiered system for each element that is based on things like daily values. We have as many as five or six levels with different set points, which would either qualify you for one, two, or three stars or no stars. The system is also designed not to unfairly favor fortification, so products that just have lots of added vitamins and minerals don't get extra credit for adding more. And it's also got a balance on the adverse side; that is, no matter how many good things you have in your product, if it's just filled with negative attributes, like if it's 55 percent sugar, the algorithm is set so that you're not going to get stars even if you have whole grains and vitamins in it if it's got too much of a negative attribute.

The system was intentionally designed to be simple: one, two, or three stars. Consumers understand it. They can make a decision if they're choosing a more nutritious product in one second. They can have their kids go out and pick the starred items.

The consumers told Hannaford, the supermarket chain in New England that originally developed that, that they didn't want food police. They didn't want skull and crossbones. They didn't want negative messages. They just wanted something to help them make a healthier choice and we felt that this approach was a really simple way to do it.

Other systems using traffic light, using multiple icons, numbers of calories, saturated fat, trans fats, sugar, salt, I believe, are very complicated. It's confusing. It's somewhat comparable to trying to read the nutrition facts label and then you have to balance it yourself: "Well, geez, it has a little more salt, but it's got a lot more sugar, so is that better than the one that has more sugar and less salt?" The algorithm we've developed takes care of that in a way that's consistent with the science.

It is discriminating, so currently 26 percent of the products. And this is where Guiding Stars is principally now: Food Lion, Hannaford, Kings, Homeland, and Sweetbay, mostly along the whole Atlantic Coast. 100 percent of fresh fruits and vegetables get stars and I will tell you we see no effect of that labeling.

People are not saying, "Oh, I didn't know fresh fruits and vegetables were more healthy, I'll buy more now." But we do see movement and I'll show that you in a minute, in the center of the store, where the processed and manufactured foods are. You can see here, 51 percent of cereals get a star, one, two or three; 50 percent of seafood; 21 percent of meat; 21 percent of dairy; 7 percent of soups and 7 percent of bakery items.

I probably don't need to tell you, but soups generally have lots of sodium. It takes them off the stars. People were very surprised. There were a lot of learning moments when we launched yogurt. "How come my yogurt doesn't get stars; it's a really healthful food?" Well, many yogurts have lots of added sugar to it. Cottage cheese has lots of salt added to it, so people were really surprised. The good news is that I'm happy to tell you Guiding Stars is a discriminating tool. The bad news is it may be a bit of a scathing indictment about how healthy the food in our supermarkets may be.

So we did look, between launching the product which over now three years ago, at the shift in products that have stars versus not. You can see that we found these star induced shifts, especially in frozen dinners, ground beef, and milk, getting people to move from full fat to our 2 percent, 1 percent, or skim milk products. Yogurt I just mentioned and ready to eat cereals is where we found the biggest effects.

If you can read the tiny print on the

percentages there, you will see we had a huge, whopping, statistically significant improvement in moving more starred nutritious food of about 2 percent. As a nutritional biochemist, I will tell you, I said, "This is nothing; this is just not important. This is never going to fly." On the other hand, my colleagues in the supermarket were deeply impressed that this was really a big change. I didn't appreciate how you're supposed to do the metrics here.

A 1.4 percent increase that we had seen over the whole store, not just the center store, translated in one month to moving three million starred items out the door. That translated into about one ton in one month of less salt going out the door, so it can have a big effect, and again this illustrates this same thing.

Now, we're looking here just at one month's sales of ready to eat cereals before launch and an after launch in about 300 stores on this one. So we saw a decrease in the number of rating cereals that had no stars and an increase in those that had stars. There's still more non starred cereals being sold. We're just seeing an increase in the starred ones and a slight decrease.

Because these are starred on the icon, the FOP had about half as much sugar, what we saw is less sugar went out the supermarket doors. More stayed behind, and this was equivalent to about 60,000 grams of sugar not being sold during that period of time. And the same is true for our algorithm, which encourages whole grain and fiber. In this case starred products have about five times as much fiber as those that don't get a star, and as a result there was about almost 19,000 grams of fiber that got sold that otherwise wouldn't have with this program, so we think it makes a difference.

Guiding Stars is now licensing its nutrition navigation system to high schools and colleges. They now have a mobile iPhone application. This is a study that was just presented at the American Dietetic Association meeting two weeks ago in a high school where in the cafeteria where breakfast was served, they looked before and after putting Guiding Stars in. As you can see, there was a slight decrease in the non-starred foods, and there was an increase overall in starred items.

So it seems to work even with high school kids, and so again I just want to tell you that this idea of Front-of-Package, at least in some circumstances, can extend from the Front-of-Package to the shelf tag, to cafeterias, to restaurants, to websites and now to mobile applications where Guiding Stars is going to allow you both to shop in advance and know you can just bring up all the starred items on your phone and so on.

So those are the comments I wanted to make to you this morning.

DR. DANIEL: Thank you.

DR. BLUMBERG: Thank you.

DR. DANIEL: Thank you, Jeffrey. If I could turn the microphone over to Dr. Leslie from Stanford.

DR. LESLIE: Thanks everybody. This is a paper I'll try to summarize in ten minutes about evaluating the impacts of the mandatory calorie posting law that came into effect in 2008 in New York City.

I want to begin by maybe getting your expectations down a little bit. There's a lot of really interesting economic questions around mandatory versus voluntary disclosure, and I know there's a number of people in the room who are probably very familiar with the unraveling hypothesis. It creates this really interesting question for policymakers around whether or not we need government intervention or whether or not there's sufficient incentives for the market to provide the level of information that consumers would need, that there are rewards for doing that.

There are lots of interesting questions around that and I think there's research that has been done and research that needs to be done to flesh out some of those issues more carefully.

In this paper, we don't really push back the frontier very much on those kinds of interesting economic issues. We're motivated to do this paper, because there's a lot of people out there who think or who say that obesity is one of the biggest policy problems facing the United States and a lot of other countries around the world today, and this seems to be perhaps the only really significant policy initiative that is specifically focusing on trying to reduce obesity.

As somebody that's done research on information disclosure previously, in fact even in the context of restaurants, although restaurant hygiene in that case, it just struck me as a really interesting policy question as to whether or we could provide any evidence that information disclosure was going to have any real traction on this big policy question, so that was the reason for doing the paper.

It's less about fleshing out the subtle economics and more about just the basic policy assessment question of, "Does this look like it's going to have any impact on obesity."

So the law came into effect in New York City in 2008. It's the first place in the world that has done this kind of thing, which is to say mandatory calorie posting on all menus for all chain restaurants. Chains are defined as any restaurant that has 15 or more units nationwide and in different parts of the world.

This law is also now being implemented, other parts of the U.S. and other parts of the country, around the world, sometimes the number is as low as 10 and sometimes as high as 20. People play around with that a little bit.

The idea is that the calorie information is supposed to be posted as prominently as price. I will show you a photo in a second of what that looks like, and some of you may know that actually this idea has been picked up by President Obama in the health reforms that were passed last March. This is going to become a federal requirement now and I think we have yet to see exactly the details of specifically how that's going to be played out, but there is in principle now a federal requirement that there's going to be a posting of calorie or nutrition information on restaurant menus.

It's unique, by the way, because that's the first time that this kind of approach has gone from packaged food. We've had mandatory nutrition labeling on packaged foods since the early 1990s, and there are people in this room that have done a lot of research on that.

It's the first time this has gone to the food that we eat in restaurants as well, which is, as you probably know, an increasing fraction of the amount of calories that people are consuming these days.

So this is a look at Starbucks and our data is from Starbucks. I'll tell you a little bit more about that in a second. So you can see that it's doubling the amount of quantitative information on the menu. In this instance calories really are being treated as prominently as price.

If you look at the way it's been implemented by other chains like McDonald's and others, sometimes you can debate whether or not it's been given exactly the same amount of treatment as price has, but that's for the regulators I guess to decide whether that's being done.

We also look at some of those numbers -- and you're probably wondering -- I would be wondering how much you would have guessed what those numbers were.

Starbucks, by the way, was often identified by journalists as exactly the kind of chain that was likely to suffer from doing this as people are shocked when they find out how many calories are in those lattes and frappuccinos and so forth that they eat.

So it became a bit of a cottage industry for journalists in New York to go and interview people in Starbucks and try to capture in a colorful way the shock and surprise they would have when they saw these kinds of numbers.

Here's the information that was already available at their website, so to be very clear, they had already had voluntary disclosure happening in this industry. It wasn't in the format that the new law required it to be. It was at their website and sometimes there would be brochures in stores as well, and other chains did different approaches.

There's one chain, Subway, which had already taken the approach of putting calories in their menus, but they're pretty unique among chains. There weren't really other chains that were doing that at all and I'm going to show you this to emphasize that some of the information was available. You had to go on the Internet to get it. You can see that there's a lot more information contained there than just calories.

There's a group of people that do research around how complicated information like this is difficult for people to understand and utilize and, something that was plausible to me at least, that this kind of information that we're seeing here on the website is not nearly as impactful perhaps as the kind of information that they have to put on their menus with calories on the menus.

Here's another example I like to show. This is the Italian chain Sbarro. The reason I like to show this one is because there's this one particular item there, the fettuccine alfredo, and that to me really highlights the degree to which this is likely to be powerful information.

It seems intuitive that you would expect that providing this sort of information would have a big impact on the behavior of consumers. On the other hand, I think it's worth noting that not only was the information already available for anybody who was interested and wanted to know that information. Some people want to know and some people don't care. But secondly, when it comes to fast food and especially for the chains like the McDonald's of the world, I think it's reasonable to suggest that people don't really care about nutrition when they choose to go and get a meal there anyway.

I've already given up on the hope of really having a nutritious meal when I make the decision to go to McDonald's in the first place. People have also shown extensively that people care more about price, taste and convenience when they're making those kinds of choices than they do about nutrition, so there's reason to believe that this might not have very much of an effect.

So we spent a lot of time trying to talk Starbucks into giving us their data and after a long process, that worked out okay. We also spent a long time arguing with them about whether or not they would let us name them, and eventually they did decide to let us name them in the study, so we have pretty amazing data for Starbucks, and I want to be really clear that that's going to be both a strength and a major limitation of the research.

We don't know anything about what's happening outside of Starbucks. Not only do we not know what's happening at other chains, I don't know what people are doing when they're eating at home or other things like that, and that would be in many ways, of course, better to have that kind of information.

The strength of this data, as I'll try to convince you, is it's remarkably deep, and so we can say a great deal about how exactly this plays itself out at Starbucks. One of the reasons we choose Starbucks is because we thought that they would probably be quite impacted by it and also because Starbucks is just so big anyway. They're the second biggest restaurant chain in the world with revenues in excess of the entire movie industry and stuff like that. So even if we don't say anything about anything else, just about Starbucks, hopefully people would care about that anyway.

So the data has three bases to it, and there's two really that I'll mention here. We have this transaction data and we have this cardholder data. All of the data goes for three months before and eleven months after. So posting started on April 1, 2008, and we have the data for all transactions in New York City, and also for control cities Boston and Philadelphia where there are no calorie postings.

Since it's 110 million transactions, it's literally every transaction at every Starbucks in every one of those locations. Then they also have this cardholder data set because the transaction data is anonymous data, and I don't know anything about the people making those transactions, although we do know the locations of those transactions. We can associate that with local demographic information, and so with some assumptions there, you can try and talk about how the effects differ by demographics.

The cardholder data is nice because we get to follow individuals over time and there are 11,000 cardholders. Actually in the data set, there are several million for the entire United States, but there are 11,000 that are engaging in a lot of transactions in New York City, Boston, or Philadelphia. That's nice because you can use individual fixed effects and see how exactly it impacts on the behavior of individuals.

This is probably the single most important figure in the paper, and I'll leave it to you guys to look at the paper if you're interested to see more details. This is basically showing you the estimated impact on calories per transaction so the vertical axis here is going to be the percentage change in calories per transaction, this log of calories per transaction.

Then the top figure is for the transaction data, and the bottom figure is for the cardholder data, so this is based on regression. This isn't just the raw data. Here we're using our controls of Boston and Philadelphia.

Here what we're doing is flexibly estimating the impact on a week by week basis, and in the dotted vertical lines is of course the moment when calorie posting comes in. We don't use that information in the estimation. We just put that on the figure afterwards, and the top figure, you can see clearly that the calories per transaction falls.

There's that jump up around the holiday period. You don't see that jump up when you look down at the cardholder data below, so we think that's mainly a compositional effect in the transaction data rather than impact by individuals. Anyway, no matter how you slice it and whatever you do with the data, you always get that there's about a 6 percent reduction in calories per transaction.

What's interesting though is all of the impact comes from food choices. The thing which is probably the most surprising, the single most surprising aspect of our study, is that people's beverage choices at Starbucks were completely unaffected by calorie posting. We tried everything we could to find evidence that people were substituting to either smaller sizes, from large to small or whatever, or to different beverages.

We also have access to their milk order data to see if people are more likely to switch to a low fat milk or whatever, and there's just absolutely no impact we can find anywhere on people's beverage choices. I'll leave it to you guys to decide whether or not that's something you would have expected.

On the one hand, you might say that people are sort of inflexible in some sense when it comes to their coffee choice at Starbucks or maybe that's what's going on there, so all of the impact is on the food. Starbucks sells less food as a result of this, and in the end, calories fall from 247 average calories per transaction to 232 and that's all driven by a reduction in food calories.

We also find the effect is greater for individuals that tend to buy lots of calories. With the cardholder data we see that. So that's good. We would hope that the people that tend to buy a lot of calories would be the most affected by this kind of thing. That's good news I guess.

However, we also find that the effect is smaller for less educated and less wealthy people, so that's kind of bad because we think that education and wealth are negatively correlated with obesity, so that says that the people that need it the most are the less responsive to this kind of information as well.

From the point of view of Starbucks, one of the nice things we can do is tell them what the impact was on their profitability. There's actually some neutral effect for them. Even though they sell less food, that's less revenue per transaction, they actually have an increase in the number of transactions per day, so it seems to be driving business towards Starbucks. So in contrast with what I think a lot of people expected that Starbucks would be hurt by this, it seems like they're actually doing okay with it.

We look to see how the impact differs for the Starbucks that have a Dunkin' Donuts right near by versus the ones that don't, and I guess it's probably obvious to people here, but I always have to tell people on the West Coast that Dunkin' Donuts is the major competitor to Starbucks.

In fact, they have a pretty significant positive impact at the Starbucks locations with a Dunkin' Donuts nearby, which is interesting I think to us because it's evidence suggestive of the fact that this kind of disclosure has differential impacts for different firms.

Some firms are going to benefit from this and some firms have a negative impact from it. Remember that Dunkin' Donuts is also having to post calories there as well.

In the paper, we talk about how large the magnitude is. It turns out to be pretty small. The 6 percent affect is unlikely, I think, to have a major impact on obesity. It will have some impact, but it's not going to be major, so we emphasize in the paper that this is not a silver bullet.

We also try to talk about how it's important to understand that the effect may be much larger at other chains. And if the effect here is mainly on food, and Starbucks is after all a beverage focused chain, then if you were to extrapolate from that, you might think that other food focus chains like McDonald's would have a much bigger impact of this, but that's clearly speculation on our part.

Lastly, something we all know already is that so much of this is about how information is being provided, so putting calories on the menu at point of purchase seems to have been pretty important. The last thing I'll mention is that from the interactions I've had with Starbucks and from what I've been paying attention to in the media and other stuff, I think it's really clear these days that there's a huge focus from all food companies to want to provide more nutritious food.

So I think the reason to feel optimistic about these kinds of things is that we're going to see some innovation around food offerings and that the supply side impacts of those kinds of things are hopefully where we're going to get the most traction and hopefully the biggest impact on obesity in the long run.

Thanks.

DR. DANIEL: Thanks, Phillip.

I can't resist a quick reaction to Phillip's paper in that there are some calorie postings in the Washington area in certain jurisdictions. I was at a Washington Nationals baseball game this summer where there are no calorie postings at the Five Guys burger chain offering there, which is a nutritionist's nightmare.

And in front of me was a man talking to his friend saying, "Thank goodness they don't post the calories here at Nationals Park; I was out at a store in Maryland, and they posted the calories, and I couldn't order what I wanted."

And so I was curious to see whether there would be any carryover affect since he knew what he was buying into, but when he got to the front he ordered a triple cheeseburger with all the fixings. I thought, "Okay, at least he knew what he was buying, and perhaps that would affect his behavior elsewhere," but the information is that the evidence was there that in this one instance didn't change it apparently. Maybe he didn't have a four bagger, I suppose, but it didn't seem to change his behavior.

I just want to take a few minutes to mention some of the disclosure initiatives. This is but a sliver of what's going on in Washington at certain federal agencies with regard to disclosure issues, some of them dealing with foods, some of them not.

Just to give you some idea of some of the initiatives, three agencies that I will touch on very quickly: My home agency, the Federal Trade Commission, an initiative going on at the Food and Drug Administration having to do with Front-of-Package labeling, and an interesting memorandum was issued this past summer by the Office of Information and Regulatory Affairs at OMB, which oversees federal regulations in a systematic way that I would just like to highlight and bring to your attention, if you don't know about it already.

The Federal Trade Commission has what are called Endorsement Guides. They were first issued in 1980, but they were updated in 2009. The disclosure issue is that if an endorser of a product has a material connection to the seller, that that material connection, meaning the financial connection to the seller, makes money from the sale of the product that consumers might not expect to be there but for the disclosure, then the disclosure has to occur.

So in 1980, when it was first introduced, some of the issues had to do with consumer testimonials about the types of reactions or experiences they had with diets or with other products, and the idea was that consumers really ought to know whether or not those endorsers have a financial reward coming their way from the sale of the product.

The FTC wanted to extend this principle to Internet commerce, and so the guides were updated and released late last year, and there's been quite a lot of discussion and interaction from the internet commerce community and the blogosphere about how best to implement what I think personally is a sensible principle.

There's been quite a lot of discussion, and I would commend for you, if you're interested in how this disclosure principle might be playing out in Internet commerce, to look at a June 2010 document issued by the Federal Trade Commission "Facts For Business: The FTC's Revised Endorsement Guides--What People Are

Asking," and that document provides the FTC's current thinking with regard to the obligations on Internet sellers as well as those who review and potentially endorse those products.

The second initiative at the Food and Drug Administration with regard to foods, not surprising, is their Front-of-Package labeling initiative. This is a different initiative from the, what I suspect is familiar to most, nutrition label that's on every packaged food product sold. This is a regulatory initiative looking at possible requirements for the Front-of-Package, the thing that consumers see while they're going through the store.

Much like Jeffrey's research and Jeffrey's interest in Guiding Stars, providing information directly to consumers at the point of sale, this initiative at Food and Drug Administration was launched following the introduction and then relatively rapid withdrawal of the Smart Choices label, which somewhat similar to the Guiding Stars program, was an industry driven program to provide incentives for manufacturers to put the Smart Choices checkmark on to their package if their nutritional profile was such that it qualified.

Like the Guiding Stars program, there were a lot of smart people who worked on the program, nutritionists as well as industry experts. When the label was introduced, however, there was quite a fire storm around the Smart Choices label because certain heavily sugared cereals qualified for the checkmark, and as a result, there was some concern that the algorithm, if you will, that underlay the determination as to whether or not that food would get that checkmark was flawed in not providing consumers with credible or accurate or useful nutrition information.

So it was withdrawn and the Food and Drug Administration then opened up a proceeding to gather information on Front-of-Package labeling. Comments were requested in April. They were due at the end of July. The Guiding Stars was among those that provided some comments to the Food and Drug Administration, hundreds of others as well.

The Institute of Medicine issued a report following that period, just about a month ago, talking about its first phase of research of what's going on with Front-of-Package labeling. That report provides information on the many, many initiatives that have been launched in the U.S. and in other countries on Front-of-Package labeling, and just very recently the GMA and the FMI, not the FLI, I apologize for the typo, the Food Marketing Institute and the Grocery Manufacturers of America, announced that they have launched their internal initiative to provide Front-of-Package labels to foods.

They didn't disclose exactly how their algorithm was going to work or what the label was going to look like, but they are seeking to be a bit ahead of the curve perhaps of what the Food and Drug Administration is doing, so a very active area.

Again the objective is to provide useful information to consumers at the point of sale, the difficulties being that perhaps a red light, green light approach that underlay is the Smart Choices was not the way to go given its costs and benefits. The more tiered approach from the Guiding Stars with one star, two stars, three stars seems to have potentially better market success anyway. It's still out there, still being used by the chains that Jeffrey identified.

Lastly, I just want to bring your attention to this June 2010 memorandum issued by OIRA entitled "Disclosure and Simplification as Regulatory Tools." It is a thoughtful document seeking to lay out the principles for disclosures in a regulatory setting, not just for the food products at issue at Food and Drug Administration, but across automobile labeling or gasoline mileage standards and that sort of thing.

They identified the seven principles for disclosures, which all federal agencies have been instructed to consider when designing their disclosure policies and remedies. I'll run through them quickly and stop at a couple with a quick comment.

Agencies should identify their goals, here seemingly obvious, but the take away that I took from that is that if the goal is a warning that this is a product you should be careful about or consume less of, then you can have a little more eye catching label, if you will, eye catching disclosure remedy in order to get that warning across to consumers more effectively. The recent cigarette pack label that Food and Drug Administration is considering, I think, fits into that goal of getting that warning to be very prominent and eye catching.

Disclosures should be simple and specific. They should be accurate and in plain language. They should be properly timed and placed. That's going to be an important and has been an important issue in Internet commerce: When should consumers see a disclosure if there's one that needs to be made at all about the relationship between certain parties active in internet commerce. Is it at the front end, at the back end when they're thinking of making a purchase? What's the right time for that to happen?

Disclosures using ratings or scales such as the Guiding Stars program should be meaningful, and then for the economists in the room, the agency should test in advance and monitor over time the effects from disclosure requirements. I think certainly I would strongly agree with that, that the agencies and private academicians, when they get a chance, should look at the costs and benefits of these disclosure regimes or disclosure proposals.

I will stop there, and we can then move to the discussion portion of this. I would like to offer my two panelists the opportunity to make comments in reaction to or in amplification of their comments made already.

Let me then turn to some questions, if I may, and I would like everyone in the room to join this discussion, which will take your lead as well as mine. We should have some microphones in the room so I would encourage you if you speak to please use those.

DR. LAIBSON: So this is a question for Phillip. David Laibson, Harvard University. As you know, there are a bunch of different research teams measuring the effect of food calorie disclosures in the

New York case, there's Brian Ebel and his collaborators. There are the New York City research groups and their imperfect reports of their results. I think they're still kind of working through that process.

My impression, I just spoke to Brian actually on Monday, in his own work, which was in health affairs awhile ago, was that they found no effect on average for a range of stores. I don't recall whether Starbucks was in their data set, but they had five to ten stores and they were using New York as a control city, and they found literally no effect. In fact, the point estimate was whether it went in the other direction, disclosure insignificantly raised the caloric consumption per transaction.

The New York City data apparently is mixed. New York City seems to be reporting selected stores that showed reductions in calorie consumption, but if you think about the aggregate of their data, it looks again like there's no effect.

So I guess my question to you is: What do you make of the fact that on average there seems to be some stores that increase caloric consumption as a consequence of disclosure and some stores that decreased caloric consumption as a consequence of disclosure? Starbucks seems to be on the decreased side.

Is it that we should be thinking about this as a kind of failure in the aggregate, or is it that somehow the Starbucks example is more indicative of where we're heading? I'm wondering what your thoughts are.

DR. LESLIE: Yeah, so one of the things they did in the Ebel study did was deliberately choose low income neighborhoods where there would be less educated people as well. That was a specific choice they made going in. We also find that the effects are smaller in those neighborhoods, and their effects are small in magnitude as well.

You simply find a positive effect, and recall a large standard error, so it's not that they're finding any real significant effect of any kind at all, and others that have looked at this as well. I think that there's a consistent theme emerging which is to the extent that there are effects, they are small, and we have so much data that we can get such precise estimates of a relatively small effect.

So I think the thing we're all on common ground with is that these are not large effects and we're sort of quibbling over whether or not it's 6 percent or zero. I don't personally believe that it's likely to be positive in any particular cases.

That seems somewhat implausible to me, but I have yet to see anybody, and the health department stuff has not come out yet, so we haven't seen that fully, but I don't think it's plausible, based on what we've seen before, that there's large effects of this going on anywhere.

DR. BLUMBERG: I would just add a comment to this in that I think time is also a factor. I think it's going to take awhile for consumers to see it and get used to it and so on. In the Guiding Stars, there was no immediate uptick when people really did have to understand it. I don't think that was the case with Starbucks.

There was a lot of advertising. There were store banners. There were shelf talkers. There were website releases. They had dieticians who took customers on store tours and a big effort was also made to educate, to train the employees of the store, even the clerks that stocked the shelves because they thought customers would say, "How come my favorite product doesn't have a star."

So we tried to give, then there were a lot of employees, at least some basic information so that people could respond to it. And then over a matter of months, and I think the big increases that I was talking about, 1 to 2 percent really were taking place after about a year.

There are other issues that we saw in Phillip's presentation, also in the supermarket. There are seasonal effects. People's shopping behavior is different around Thanksgiving than it is in June, that sort of thing, so you need time I think to evaluate these things.

DR. DANIEL: Yes. The gentleman to my left. My apologies. The microphone is going to do the talking. We'll go this way, and then we'll get to you next.

DR. ZINMAN: John Zinman, Dartmouth College, A question for Tim: So big inputs to any cost benefit calculation in this domain relate to enforcement, enforcement costs and enforcement effectiveness, so could you just shed some light on how that FTC is planning to deal with enforcement issues on these new initiatives?

I know in the past the FTC has often been charged with enforcing mandated disclosure in various markets, but not always been allocated the resources needed to enforce them effectively so it seems like an important issue.

DR. DANIEL: Two things that I'll say. One, I agree that it's an important issue. Two, I want to make the disclosure that every FTC employee makes, which is that I don't speak for any Commissioner. I don't speak for the Commission. These are my comments only.

I would say in terms of enforcement, the FTC certainly has enacted, I think, a conscious and pronounced information effort with regard to getting information out on new initiatives. If it is the Endorsement Guides for instance and their impact on Internet commerce, speeches are given. Conferences are attended. Disclosures are made via these things like facts for business that go out to the industry and to those that are participating.

In terms of how they're going to enforce these, I'm really not in a position to tell you being a relatively new economist. The economist matters as to whether I can tell you what the FTC is going to do in a law enforcement approach.

But they did bring a case in the past, settled a case in the not too distant passed called Reverb, which was a firm that was alleged to have put disclosures on to the iTunes site or endorsements on to the iTunes site that were written by their employees, and therefore not everyday consumers, and that was deemed to be a deceptive act or practice.

So the FTC is watching what's going on in Internet commerce. I wish I had a full answer for you, but that's kind of where we're going.

MR. JEITSCHKO: I'm Thomas Jeitschko from the Department of Justice. I have a question to follow-up on the earlier question for Phillip.

I understand why, from the perspective of studying the issue your benchmark is calories per transaction, but of course from a health perspective, that's not relevant unless the number of transactions stays constant.

Do you know if any of the studies have tried to look at what actually happens to caloric intake across the board? In particular, you're worries about if I switch from a doughnut to a bagel in the morning or something like that, but then it turns out at ten o'clock, I add another Snickers bar that I wasn't used to eating before, and it might not do all that much good.

Do you have any insight in that?

DR. LESLIE: Like I said at the beginning, unfortunately one of the big limitations of the data is we're not able to say anything about people's caloric purchases outside of Starbucks. The one thing I can say that speaks a little bit, I think, to your question is that when we look at total sales of calories at Starbucks, which combines of course calories per transaction, total transactions, that actually ends up going down by a small amount, less than 6 percent.

We see 6 percent reduction in calories per transaction. There's a slight increase in number of transactions, so the total decrease in number of calories sold by Starbucks on average per day, I think the number was around 4 percent.

That's not the same thing I know as you were asking, but we don't have any data on individuals, and actually I think one thing we probably need to do better is tracking, getting that data and making that data available to researchers.

DR. BLUMBERG: I would just add that when the Guiding Stars program was launched in Hannaford's, they had no loyalty cards, no frequent shopper cards, so there was no way of tracking the individuals. It was just number of items sold, but that Guiding Stars program, and there are other Front-of-Package programs, are now in chains that have those cards that actually track everything the consumer buys in that store.

If you in fact shop at two different supermarkets, then obviously there's no way to track what everybody purchases.

DR. DANIEL: If I could ask Jeffrey, he and I spoke briefly before the session began, and he mentioned quickly in passing that some lessons were learned along the way by the Guiding Stars panel, and this is obviously a mix of industry and government or industry and private and academic efforts, just what a couple of those lessons were and what surprises you may have encountered along the way. I would be curious to hear that.

DR. BLUMBERG: I have to recall our conversation earlier.

DR. DANIEL: You can start from scratch.

DR. BLUMBERG: I am a nutritional scientist. I do a little bit about foods and nutrients, but I really knew nothing about retail sales in supermarkets, but when I was originally invited and the supermarket chain said, "We hear our consumers are confused about making food choices, and we think if we institute some system, we will be better able to retain our customers and maybe even have them spend more of their shopping dollars at our store, and maybe even attract some from competitors."

So it was a noble effort to educate customers or give them information about which are your more nutritious choices, but the business model was to try to retain and grow their customer base, and I said, "Sure. I mean, I know a lot about food. I can tell you what's a good choice or a bad choice. This is not going to be hard."

It was really difficult. It's extraordinarily complex when you try to take 60,000 items and come up with a universal algorithm that rates them all fairly and across the board. It took over two years and more than two dozen iterations of refining the algorithm until it was both consistent with established science and made sense in the end.

There were lots of things to struggle with. One of them is that I think virtually all of the programs -- I can't speak for Smart Choices because that was an industry program, so a company that put a label on their box really knows what's in there -- but those that are based on supermarkets relying on the nutrition facts label. And the nutrition facts label has some good information and it has absolutely no information about many of the other things that we were looking at.

A food label, for example, gives you the total sugars. It doesn't say how much came in the natural food and how much was added just to sweeten the product. There's a lot of information that would have been nice to include in an algorithm to do these kinds of evaluations, but the information simply isn't available, and no supermarket chain can go to every manufacturer of every product and get that information, which is proprietary to start with.

Tim did mention-- again I have my biases -- the utter fiasco of the Smart Choices program. A lot of effort went into that and the industry worked with the American Dietetic Association, the American Society of Nutrition to get lots of endorsements and present themselves as really credible. But I think, in part because it came directly from the food industry, one of the things they did was they developed 19 different algorithms, so we just wanted to make sure that within each category a formula was used.

So breakfast cereals were treated equally, but that had no relationship whatever to bakery items or to dairy items or to vegetable items, and so you end up using that kind of segmented algorithm being able to take a breakfast cereal that had 55 percent sugar, and yet still give it a checkmark.

I would tell you in the Guiding Stars

Program, we really struggled to have a single universal algorithm. It was impossible. You just cannot say meat isn't as good because it doesn't have whole grains in it. It's a different kind of food, and you also can't use the same guidelines for infant formulas or foods for toddlers. There are different requirements.

So we ended up with a general grocery algorithm, one for meat, poultry, dairy, nuts. These are all among other things naturally high fat foods. If you buy vegetable oil, it's 100 percent fat, so we had to look at ways to make the evaluations fair, but we were even unhappy having to have four different algorithms across 60,000 products.

So it was a huge challenge, and so too I hear for the FDA and the CDC, which have commissioned the Institute of Medicine to sort of look into this. I don't think I need to actually warn them at all, but it's not an easy issue. You just can't readily define good food, bad food.

DR. DANIEL: Okay. Question in the back? Thank you.

MR. RAHKOVSKY: My name is Mr. Rahkovsky, I have two questions; one is for Phillip. I have a question about not a thought in change of calories, but the change in choices, so do you see -- because information basically doubled, a lot of the information.

Do you see less experimentation of consumers so they make more choices like they did in the past? Have you looked at this issue at all?

My second question is to Jeffrey, so you look at the good or bad nutrition, but a lot of research on obesity says portion controls and energy density are two important issues. Do you think there is anyway that these things can be incorporated in the star system whatsoever, and what's the best way? what's your take on that? Thank you.

DR. LESLIE: We didn't look very explicitly at this experimentation issue. With the cardholder data, we can look and see how people or how individuals have changed their choices over time, and we did look at that very closely. We essentially found no changes certainly in their beverage choices.

We did not look at experimentation type stuff, although my guess would be that the people who have Starbucks cards are people who are very loyal Starbucks customers that know exactly what it is that they like at Starbucks, so it's probably not a really good group of people to be looking for that kind of experimentation.

Now, in terms of how people change their choices of products more generally, on the food side where all the impact is, it's worth noting, and I think this is somewhat interesting, that Starbucks had very few food items on offer that were less than 400 calories. They're not actually a lot that they advertise, but they don't sell any food items that are more than 500 calories in terms of their nationwide offering.

Sometimes there are regional offerings that are a little bit more than that, and for whatever reason, they don't advertise that, but what that meant is that if you look in the data, they have a lot of foods that are all in the 400 to 500 calorie range, so they didn't actually have many chances for people to substitute to a 250 or 200 calorie item.

When I said that calories went down by 6 percent, it was all on the food side, and three quarters of that was driven by people substituting to not buying a food item, and a quarter of that is driven by people substituting to a lower calorie food item.

So one of the things that Starbucks has done in response to that is worked hard to create a broader line of calorie offerings, so they now have more in the 200 to 300 range. That's also interesting because it highlights how one of the really interesting aspects of disclosure can be not just the demand side effects, but the supply side effects and how it causes the firms themselves to change the product offerings. I think we need to do more research around that piece of it as well.

DR. BLUMBERG: And to answer the other question about nutrient density and portion control and all that, those are really important things. It goes beyond what I think any Front-of-Pack point of purchase thing can do. I'll give you one example where I, as a nutrition scientist, have sort of mixed feelings.

The biggest shift we saw were people going from no star to one star items. We didn't see big shifts going up to three stars, fresh fruits and vegetables, for example. So we saw movement, people bought fewer full salt, full fat potato chips, and they bought baked, low salt potato chips. And I'm thinking, "No, I really wanted them to buy carrots and broccoli, and that's not how it worked."

I mean, they did make a less bad choice as a result of the icons. It's not exactly moving to follow the dietary guidelines for Americans, however. I do think we need so much more research to evaluate this. I think there's a real opportunity in using those stores that have the loyalty cards that give you discounts and stuff because it really tracks all the food items, and they know -- I'm not sure how much -- who you are, so they can get much more individualized data.

DR. DANIEL: Looking to my panel. I have one more question.

DR. PAPPALARDO: I'm Jan Pappalardo, FTC. This is really interesting research. It's really important. What we regulate at the FTC on the consumer side is information for the most part.

The one thing I was wondering about was if any of you had considered looking not only on the change of behavior but on the change in knowledge. There's a sort of philosophical question of what is the objective of disclosure. Is the objective to improve consumer understanding and knowledge or is it to change behavior?

It's quite possible that if you observe a change in behavior, it may be because people did understand, or maybe they misunderstood. We found in our own research on disclosures that it's very, very important to understand how people interpret the information.

So, for example, if you don't see the desired change in the behavior, is it because people misunderstood the disclosure because it was not tested on real consumers? So I'm just wondering if you have any plans to do a holistic project where you look at all these stages.

DR. BLUMBERG: Who is that addressed to? I can tell you that the Guiding Stars licensing corporation is interested in supporting research. I think they're more interested in licensing products, but to the question about information versus changing behavior, I would say our primary goal is changing behavior.

A star tells you nothing. We try to educate people about what the star is based on, but we're trying to get people to just make more nutrient dense, nutritious choices because they see the star. We actually found out, at least when the program was launched, that people were doing what I like to see, they looked to find out how come this has two stars.

They turn it over to look at the nutrition facts label to see why it was better, but by and large what we're finding from asking the consumers is, "We just want something that's really simple, really fast, really easy," so they can get in and out of the supermarket.

DR. PAPPALARDO: A follow-up comment: We've done research on appliance labeling and other things where you look at stars, and people sometimes misunderstand the star to think that it means something beyond what the dimension is that the star relates to, so to try to make sure people are really making decisions that they fully understand which seems to be a great place to begin.

DR. BLUMBERG: I think that's a really interesting point, but no, we don't have that information.

DR. DANIEL: We are going to need to wrap it up. Phillip, I know you had lots to say on that question. You can maybe do it in the coffee hour.

We would like to thank our panelists for really a terrific discussion and very interesting research on a very interesting area and I will turn it back over to Pauline.

(Applause.)

DR. IPPOLITO: Well, it's my great pleasure to introduce our keynote speaker for the morning.

Roman Inderst is a chaired professor at the University of Goethe at Frankfurt, and before returning to Germany, he was at the London School for a number of years.

He writes in a number of areas, has fraud interests, but particular corporate finance and banking, competition policy, and information economics. We've been following his work with great interest because he's moved towards retail finance in recent years, and that's obviously been a big concern of ours, as you may have known from reading the papers.

He also has been looking particularly at payment schemes in markets, how is a broker compensated and how does that affect both market performance and the quality of expert advice that the broker gives customers.

Finally, let me just mention that this year Professor Inderst has been awarded both the Leibnizpreis prize for scientific achievement and the Goshen prize, which is the German award for the best economist under the age of 45, so it is my great pleasure to introduce Roman Inderst.

DR. INDERST: Thank you very much for the opportunity to talk here today, and sorry for the delay. The talk today I think is very focused on financial advice. I think it is nevertheless interesting. I get I hope all the other areas because it deals with issues of disclosure, and we've just talked about disclosure in the food industry, and it deals with issues of advice, and of course advice is also important in our industry as well.

For those points here, I wanted to draw your attention to the first two papers. There's a paper with Mark Ottaviani, which came out in Competition Policy International which takes a very generic issue of advice and from a non technical perspective.

Next week, precisely on Monday, there's going to be a presentation in Brussels, and Dave is going to be there as well, of a report that we did for the European Consumer Protection Agency on consumer decision making in retail investment services, so there is a report and experiments, et cetera, and in that talk today, we will go over a little bit of this as well.

I would like for the next 25 minutes or so is to first just revisit some of the key features of the market of financial services, maybe why we should think about consumer protection in the first place, and that's going to help me to zoom in on professional advice and it's important, so it's better to make my case then, and then talking about the shortcomings of professional advice, of a couple of empirical and particular theoretical studies, and finally draw policy conclusions, if time permits.

First, what do we know about advice and its role in the market for retail financial services?

Consumers' decision-making problem is extremely complex. They must decide not only when and how much to save, but also in which asset classes to invest and, ultimately, which individual assets to buy, and then how to monitor and readjust their past investment choices. Their choice is made more complicated by the sheer number of financial products.

Consumers often lack the most basic financial knowledge, including knowledge of standard products or of standard concepts such as inflation. In addition, they may lack the cognitive and numeric skills to perform the most basic financial calculations.

Finally, consumers may have deep-rooted problems when it comes to dealing with risk, in particular financial risk, and when it comes to decisions that are made for the long term. Then, some of the heuristics that they use in their everyday life may perform badly, or consumers may excessively shy away from options that they perceive as being too uncertain or too ambiguous.

Altogether, our findings show that there is potentially a large role that advice can play to enhance consumer decision-making, for instance by reducing complexity, by providing information or by educating consumers about potential misperceptions that they may entertain.

A web-based data collection method was used in eight European Union Member States to survey 6,000 consumers, half of whom had purchased retail investment services within the last five years [see Chapter 4 of, "Consumer Decision-Making in Retail Investment Services: A Behavioral Economics Perspective," Final Report, November 2010]. Our survey showed that consumers indeed turn to financial advice or receive it when they make investments. It seems remarkable that even though our survey was online, i.e., among consumers who are "internet literate", still only fourteen percent of purchasers of financial products said that they had no contact with advisors. This finding is in line with other national surveys.

Still, this does not tell us much about the impact of advice. What if consumers obtain advice but ultimately do not or need not rely on it? What if advice was just a somewhat unnecessary by-product of the role of brokers or other intermediaries as facilitators of financial transactions?

We think that our research is valuable also as it speaks to the real importance of financial advice, and thereby suggests the importance of this area for policy making. When setting up our purchase process review, we took care to sample both recent purchasers of financial products and non-purchasers. And, indeed, there is a remarkable difference in attitudes between these two groups. The fraction of purchasers who trust financial advisors is 60 per cent larger than the respective fraction of non-purchasers. Purchasers are also much more likely to trust other financial institutions or intermediaries.

The same picture arises when we look at other questions. Purchasers are much less likely to believe that financial institutions and intermediaries suggest products that are unsuitable just to make a sale. And purchasers are also much more likely to think that their advisors have the necessary expertise and knowledge.

When we put this data, together with a number of controls, in a simple, preliminary regression, this suggests that, all else equal, a customer who trusts financial advice is 12 per cent more likely to purchase a financial product. Financial advice and the trust in it thus seem to be indeed important determinants of consumers' decision making.

Our data allows us to go one step further. We can ask for which products advice is more important. Intuitively, these should be financial products and services where consumers have less trust in their own capability, for instance as these are particularly risky or unfamiliar. Let us pick investment in the stock market.

On average, 17 per cent of consumers in our survey say that they have purchased stock or that they are already invested in stock. Investing in stock may seem particularly challenging for consumers who lack the necessary skills and knowledge - that is, unless they can rely on financial advice to bridge this gap!

A first, simple regression suggests that trust in financial advice has indeed a particularly strong effect on the propensity of less educated households to hold stock. In contrast, this is much less the case for more educated households.

For more educated households, instead, it is their perception of consumer rights, rather than their trust in financial advice, that has the most significant impact on their decision to participate in risky financial assets such as stock. This is intuitive, as more educated and better informed households may need to rely less on advice.

In terms of policy, this is an important insight. Ensuring that the market for financial advice works and that it is trusted by consumers may benefit, in particular, households that are less capable to make their own decisions. Instead, increasing consumer protection more generally, or rather the perception of it, may affect more the decision of self-reliant, possibly more educated households.

These policy recommendations are also supported by a complementary study that I did using Eurobarometer data [see "Financial Advice and Stock Market Participation," July 2010, with Dimitris Georgarakos]. It supports the picture of different consumer segments that rely to different extent on financial advice and for which, consequently, policy that is directed towards advice has different relevance.

My comments so far point to the potential positive role that financial advice can play. One "bright side" of financial advice is that it can help to create a more level "playing field" among different consumers. But recent contributions from our literature survey also point to a "dark side".

There is growing evidence also in the academic literature that advice may induce greater churning of assets and may steer consumers towards products with higher fees and higher commissions. For instance, one recent study that I did with German data shows that, controlling for many factors, consumers' reliance on a bank's advice was the most important determinant of their security trading, affecting not only how often they buy and sell securities, but also their choice of assets.

When they act on recommendations and advice, consumers may not be sufficiently wary about a potential conflict of interest. What does our research have to say on this?

In principle, such a conflict of interest could express itself in two ways. Following a recommendation, consumers could undertake a transaction with a bank or a broker, even though it would have been better to turn somewhere else. Alternatively, the advisor may steer consumers towards particular transactions.

The first problem seems more likely when consumers purchase products through a provider's own staff. Our survey indicates that this is the case more than half of the time.

In addition, consumers seem to be largely unaware of the inducements that product providers pay and how these are passed on, both through commissions and through implicit incentive schemes. The vast majority of consumers either think that the advisor or sales person through whom they made a transaction is completely unaffected by such incentives or report that, at the time of purchase, they did not give this any thought.

This is particularly prevalent in the case of a product provider's own staff. There, 51% of all surveyed purchasers of retail investment products thought that the respective employee was not influenced by incentives at all and 36% reported that they did not think about this when making a purchase.

For other sales channels the picture is not too different. Again, the overwhelming majority of consumers either do not think that incentives influence recommendations or they ignore this issue altogether.

Based on casual evidence as well as our own field studies, we believe that consumers both vastly underestimate the importance of such incentives and pay too little attention to it. Put differently, they seem to be inattentive to the fact that their advisor or salesperson could be biased.

In fact, when we asked them directly about whether they think that, for instance, a product provider's own staff receives contingent remuneration, more than one third said that they did not think so. This is not surprising, given that less than one third reported that they saw written information or were told so verbally.

In our online study we also undertook fully incentivized experiments, and we supplemented this with laboratory experiments using almost 500 subjects in three countries. The benefits of experimental studies is the ability to test hypotheses in a controlled environment. In our context, this includes isolating different aspects of advice [see Chapter 7 of, "Consumer Decision-Making in Retail Investment Services: A Behavioral Economics Perspective," Final Report, November 201].

A situation of advice often incorporates an element of trust: Based on his better or more specific knowledge, an advisor makes a recommendation whose merit cannot be fully verified by the advisee. We analyze this aspect of advice in a setting of "cheap talk". In this setting, the advisor has only privately observed information about the suitability or profitability of a highly stylized investment opportunity. Based on this information, he can make a recommendation.

At other times, the main task of an advisor may be to produce information that the advisee can then separately verify so as to judge whether a particular product suits his needs. Still, an advisor may be able to act strategically by selecting which information to show, thereby putting emphasis on some aspects but possibly withholding other information. An advisee must now solve the slightly more complex task of putting the information that he receives into this context: To what extent does he believe that it was given to him strategically?

Finally, even when an advisor is not better informed about particular product characteristics let alone the needs and preferences of a particular customer, he may still be able to influence the customer's decision. When talking to the customer, he may stress some facts or try to play down other. Also, he may try to talk the consumer out of some perceptions or misperceptions that the consumer may entertain. We also analyze such a setting.

Though these different studies span a wide range of advice situations, some caveats should be noted.

Clearly, advisors and sales people also serve different functions, such as simply that of facilitating transactions. Also, our set-up should in no way convey the picture that an advisor can be of any value in predicting market movements, helping to pick particular stock, or providing any other "tips" of this sort. In fact, this seems to be a common misperception from which retail costumers suffer. Finally, all our experiments abstract from the dynamics that could arise when advisors and advisees repeatedly interact.

In the experiments, we varied advisors' incentives. In the baseline case, advisors were paid merely for participating in the experiment. They were put on a flat wage. In other treatments we put advisors on an incentive scheme, making their remuneration dependent on the advisee's subsequent choice.

Precisely, when the advisee had to choose how much to invest in a particular, stylized investment opportunity, then the advisor was paid more the higher was the advisee's investment. For robustness we allowed the advisor's commission to take on various forms. I will, in what follows, not dwell on these details, but must refer you to the full report. I only would like to mention that also the advisor's recommendations were generated in the experiment.

In all our experiments the advisor's incentives were disclosed to advisees. I take now first the case of our online experiments. In the baseline treatment, disclosure was provided in a neutral way and in the same font size as all the other information that subjects received.

Instead, in a second treatment advisees were given, in addition, a "health warning". This read as follows: "Notice that this means that the advisor did not necessarily have your own investment earnings in mind when he gave his advice." In a third treatment this warning was given in large red front.

Our research questions in this experiment were straightforward. How do advisees react to different recommendations of the advisor, and how does this depend on the advisors' disclosed incentives? And how does their reaction depend on personal characteristics?

In what follows, I confine myself to one particularly stark result. In the online setting subjects' reaction to a disclosed conflict of interest was extremely weak. Recall that in this setting advisees have to blindly trust this recommendation, as they cannot verify the information that the advisor privately obtained.

It made very little difference whether subjects were told that their advisor was on a fixed wage or whether he was paid proportional to what they invested. Providing a mild "health warning" had some impact, and providing a strong "health warning" shaved another bit of the average investment. But altogether these responses were very modest.

The outcome was even more extreme in the "strategic disclosure" setting. I only want you to recall from what I said beforehand that in this setting it is slightly more complicated for advisees to figure out how an advisor can try to influence their beliefs, namely by withholding bad and only showing good information. With this additional complication, however, even a strong "health warning" had no impact at all.

Let me first dispel the possibility that we do not see much in the data as they are pure noise. Note first that subjects did not blindly respond to the questions. Indeed, they strongly react to advice. There are large difference in the investment, depending on whether it was done without advice, after a positive recommendation or after a negative recommendation. Also, we see in the data that subjects that are less risk averse invest significantly more. Finally, when we control for the time that subjects took for the whole study we see a slightly larger effect of disclosure, as those who spend more time reacted somewhat more. But once again the effect remained very small.

Before I can draw some tentative conclusions from these observations, we have to look at the outcome of the laboratory experiments.

We had to keep down the number of possible treatments that we looked into. For this reason, the first key difference of the lab experiments was that we provided advisees with full and precise information about the advisor's compensation, depending on the advisee's choice. Also, the advisee was matched in each advice situation to a particular advisor, albeit anonymously. In addition, in contrast to the online experiment, in the lab subjects were in a highly controlled environment, without distraction, and with much time at hand to think through the implications of their choices.

Taken together, these differences between the online and the lab experiments account for the stark difference in the outcome.

Given the incentives that advisors received and their actual behavior, in the laboratory advisees reacted strongly to a disclosed conflict of interest. In fact, in some settings, which I explore next, advisees exhibited even a contrarian reaction to their advisor's recommendation, much like a knee-jerk reaction, even though this was not justified. But this knee-jerk reaction was much mitigated by communication, which was another feature that we were able to analyze in the laboratory.

In half of the treatments that we conducted in the laboratory we allowed advisors to communicate with advisees. This was not done face-to-face, as this would have opened up channels of communication that we cannot control. I come back to this later.

Over a certain amount of time, subjects were free to communicate via their keyboards. An immediate benefit from this set-up is that it allows us, still in a very stylized and thus controlled setting, to see whether adding some more realism in the form of two-way communication makes a difference.

Surprisingly, communication sometimes makes a difference. This is surprising as communication is not face-to-face and is conducted only through a keyboard. In addition, our decision problems are highly stylized so that, simply put, there is not really much to communicate.

So what difference does communication make? Take the setting where we analyze strategic disclosure. Recall once again that then the advisor is able to affect the advisee's beliefs by withholding bad information and revealing information only when it is good. Thus, when an advisee expects the advisor to behave strategically in this way, he should adjust his beliefs accordingly.

This is indeed what we find when we do not allow for communication. Recall that in the laboratory advisees strongly reacted to the disclosed incentives of the advisor. When it was disclosed that the advisor had biased incentives, an advisee reduced his investment by one third. But when we allowed for free communication, then this reduction was smaller. While we have not yet analyzed the protocols, we must suspect that communication allowed biased advisors to make subjects less wary, thereby mitigating the effect of disclosure of incentives. Direct interaction and communication may thus, even in such a stylized environment, undo the implications of disclosure.

Recall now that we also analyze a setting where the advisor has no privileged information. All that is relevant is also known to the advisee. And the choice setting is still very simple. So what is then the potential role of the advisor?

This experiment used choice problems from the online experiment. These problems were again framed as more realistic investment decisions, with all the advantages and disadvantages that this brings. The more we now give up control over subjects' framing and beliefs, the more there is also scope that their prior perceptions or misperceptions influence their decision-making, or at least their preferred choice without advice. Also, subjects may become increasingly insecure, in particular when financial decisions are less familiar to them. Altogether, this creates the possibility that advice by a second but not better informed person, our advisor, can influence the decision. And indeed it does so.

Given the time constraint, I only want to bring out one result. We find that without communication there is a knee-jerk reaction to an adviser's recommendation. When the advisee knows that the advisor benefits more from one option than from the other and when this option is recommended, then he is more likely to decide against this recommendation. Such a knee-jerk, contrarian reaction was also observed by disclosure experiments that the Federal Trade Commission conducted.

At first this suggests a drawback of a disclosed conflict of interest. It may undermine the benefit of advice, making the advisee suspicious, even though he need not be so in the particular circumstances. However, we find that communication allows the advisor to mitigate this knee-jerk reaction. Hence, studying disclosure only in an environment without communication would have generated potentially misleading policy advice.

Further, our experiments of persuasion suggest that while advisors can use communication to their advantage and even reinforce advisees' perceptions or misperceptions, communication is much less powerful in talking advisees out of such misperceptions. Our analysis in this regard is, however, still preliminary.

Our studies on advice lead us to the following tentative policy recommendations. Financial advice should be a priority for consumer protection in the area of retail investment services. It plays a critical role, and in particular so for consumers who are less financially capable on their own.

Further, policy makers should not take it for granted that consumers are sufficiently wary of the potential conflicts of interest in the market. This holds when it comes to the commissions and other inducements that advisors and sales people receive. But it also seems to apply with respect to consumers' perceptions of how "tied" advisors and sales people, in particular product providers' own staff, are incentivized, both explicitly and implicitly.

With respect to disclosure, policy should not put too much faith in the unspecific warnings of a conflict of interest. Our results suggest that disclosure has not only to be clearly visible, but it also has to be specific, potentially detailing the precise value of an advisor's or sales person's contingent payment. We are aware, however, that such disclosure could tilt the market in favour of vertically integrated providers, where incentives may be given implicit, say through promotions or wage increases. This must be born in mind, as the last thing that one would want is to hamper the best ally of consumers, which is competition.

Further, disclosing conflicts of interest may not be enough and it may also not be a policy to be conducted in isolation. In particular, we find that the effects of disclosure may be mitigated even by the communication that non-professional students conduct anonymously over keyboards.

Our study did not look into the various levers of influence activity that professional sales people have at their disposal in face-to-face situations. However, it must also be said that, on the other hand, our study did not take into account various forces that could induce a more benevolent behaviour by advisors in the real world, such as self-imposed codes of best practice, reputational concerns or, in particular, supervision and the threat of legal action.

"Mystery shopping" exercises are frequently performed by supervisory authorities and they have also been undertaken by academics, such as already years ago by Thorsten Hens from Zurich. Collecting data from such exercises could represent a valuable next step, albeit certainly restricted to certain providers and products.

Finally, as long as advisors' incentives remain biased or at least frequently not in line with the interests of consumers, policy makers should not count on professional financial advice as a cure for consumer misperceptions or other so-called "biases", as we discussed them earlier.

To align the interests of advisors and advisees, it could be envisaged to ban commissions or similar inducements. A policy in this spirit seems to be pursued by the UK's Financial Service Authority.

Clearly, before undertaking such drastic steps we must be clear about what is the market failure that we seek to address thereby. If the current practice of how consumers pay for advice, namely mostly indirectly through higher prices and commissions, is grossly inefficient, why is it so prevalent?

Our study points to one possible source of such a market failure: Consumers' ignorance or naivite with respect to advisors' incentives. As I have shown in work with my co-author Marco Ottaviani from Kellogg, such naivite leads indeed to an equilibrium outcome where advice is not paid for directly and thus remains biased.

However, a policy of banning commissions should clearly not be chosen before we cannot rule out safely that it would lead itself to serious inefficiencies. This could be the case when it would induce consumers to excessively shy away from advice.

Our online experiments yield some first, albeit very preliminary data on this issue. We tested subjects' willingness to pay for "hard" information, that is information that they saw directly on the screen.

We were interested in situations where consumers are uncertain whether to invest at all or whether to switch out of existing assets.

In this case, when they pay for advice through commissions, they only pay when they make actual use of the thereby gained information. Instead, when they have to pay up-front for information, this payment is unconditional. It is incurred regardless of their subsequent choice and thus, from consumers' perspective, represents a sure "loss".

Some findings in the psychology literature suggest that consumers may be excessively loss averse. I must refer you to the report for details. Our preliminary findings suggest that almost one third of subjects showed behaviour that is at least consistent with such loss aversion. As we have, however, seen in other experiments, this may depend much on the online context, which is why we would recommend to undertake further studies in this direction. Likewise, it would be extremely useful to make use of data collected in the industry or even from field experiments.

In terms of policy advice, our suggestions aim at supporting consumer decision making. To the extent that this is viable given the complexity of financial decisions, simplification and standardization of information seems a key priority, given the results from our experiments.

But this may not be sufficient, as a large segment of consumers must rely on professional financial advice. In this respect, our study focused mainly on how to address potential conflicts of interest, in particular through appropriate disclosure. We did not investigate other policies such as minimum qualification requirements or the imposition of stricter fiduciary duties and tighter supervision.

Clearly, any of the policy recommendations that we summarize on this final slide must be considered in light of all other policies that are chosen with the aim of improving the market for retail financial services in the interest of consumers.

DR. ROTHSTEIN: We'll dispense with questions, but thank you very much. We will reconvene at five after the hour, please, five after 11:00.

(A brief recess was taken.)

PAPER SESSION ONE: TOPICS IN EMPIRICAL IO

AVID NEVO, Northwestern University
PRESENTER: MATTHEW GENTZKOW, University of Chicago, Booth School of Business
DISCUSSANT: MATTHEW WEINBERG, Bryn Mawr College
PRESENTER: KATE HO, Columbia University
DISCUSSANT: KEITH BRAND, FTC
PRESENTER: NATHAN MILLER, DOJ
DISCUSSANT: ALLAN COLLARD-WEXLER, New York University

DR. NEVO: Let's settle down and get going and try to get this thing back on time with the next session.

DR. ADAMS: I'm Chris Adams. I'm an economist here, a staff economist here at the FTC. The next session is going to be chaired by Aviv Nevo of Northwestern on various topics of empirical industrial organization.

DR. NEVO: Thank you. So we have three great papers, and we're going to start with the order of the program with Matt going first. We'll have each paper, and I guess we have 20 minutes on the paper and seven minutes on the discussion, exactly seven minutes, not six and a half, not seven and a half, exactly seven minutes.

What I propose that we do all three papers and then open for questions at the end on all three papers, whatever time we have left, so Matt?

DR. GENTZKOW: All right. Thank you very much, so this is joint work with Bart Bronnenberg and J.P. Dube, and broadly the motivation for this paper was the observation that consumers frequently are willing to pay a lot to buy particular brands. This is true even in situations in categories where the physical products involved you would think are pretty similar, and nevertheless, for Coke versus Pepsi or for this kind of beer versus that kind of beer, consumers' willingness to pay is high and the observation that that has a big impact on market structure and on firm behavior.

In particular the idea that if there are things that firms can do that impact the formation of these preferences by consumers in a long term way, that would be very valuable to firms in a sense that firms in the real world in fact do spend a lot of time thinking about this.

So we're trying to understand, on the consumer side, something about the origin and evolution of brand preferences. There's a lot of theory that speaks to where those preferences might be coming from. They're models of habit formation, models of learning

-- I like Toyota cars because I've had some experience with Toyota cars and knowing that they're good, and therefore I don't know about GM cars -- models of advertising, models of social influence or social learning. And people have observed for a long time that these kind of preferences could be a really good source of market power and economic rents.

So in Joe Bane's original work on entry barriers, he sort of speculated that these preferences for branded product could be the most important of all barriers to entry in markets.

On the empirical side, I think it's fair to say we have relatively little evidence about where these things come from, especially little evidence over long horizons, so there are literatures trying to estimate advertising effects. Those, to the extent that they find any effects, the consensus is sort of the effect of my ads today is gone over a horizon of three or four or five months, which makes it hard to tie that to anything like these long-term preferences that the firms seem interested in, similarly for estimates of habit formation and switching costs.

What we are going to do in this paper is jump off from an observation that my coauthors made in an earlier paper, which is if you look at supermarket type products, for example, canned coffee, Folgers and Maxwell House, categories where the products physically are extremely similar, but the relative market shares of different brands vary dramatically across space in the U.S. This paper shows big circles of where Folgers is popular in the west, Maxwell House is popular in the east and the Midwest.

They argue in this paper that it was largely to do with who entered these markets first a hundred years ago, even though in all of these markets you can now buy both of these product. So, that's a very nice paper. That was their work.

What we're going to do in this paper is to tie this to a survey that we did of about 48,000 households who are in the Nielsen Homescan panel, which means we know all of their supermarket purchases where we ask them basically, Where were you born, how long did you live there, how long have you lived where you live now.

Well, let's look at two consumers who currently live in the same place, and they therefore face the same availability of products. They face the same prices. They face the same advertising. They face the same promotion, but they differ in where they lived in the past, and importantly, where they lived many years ago. Say, two people who have lived in Washington, D.C., 50 years ago, but one of them 50, 51 years ago was living in California and the other in Washington.

At the end of the paper, I probably won't have time to talk about this much, but we kind of used that data to estimate a model in which consumers have a stock, have a preference capital for these brands, which in the model is a function of past consumption, and we use that to think about the implications for things like first mover advantage and market power.

So to preview what we find, there are big effects of past experiences on current purchases. That explains about 40 percent of that cross state variation that they observed in the other paper. The remaining 60 percent, we can't say exactly what it is, but we interpret that as being about differences in things like prices, shelf space, availability, advertising, those things, and you can see that those things are correlated with those market shares in a way consistent with that.

The second core result is these preferences are extremely persistent, and so the advertising literature tends to find things dissipate over six months. We find even after 50 years, there's a measurable significant impact of my consumption back then on my behavior today, and then in the context of the model, we estimate that implies big barriers to entry.

If you get to enter the market five years before I do and therefore you can accumulate the stock of capital among consumers. In order to overcome that, I would have to discount the price by a lot for a long time to get back to parity, and we also look at -- we don't say much about causality in this part.

But we look at how does the importance of this kind of brand capital differ across categories, across different types of products in the supermarket, and show basically it's more important for categories that are highly advertised. It's more important for types of products that are socially visible, things like soda and beer that you eat and drink and consume with friends in a context where they can actually see what the brand is.

We have data from the Nielsen Homescan panel. This is data that lots of people use. You buy things at the supermarket, and when you get home, you scan the bar codes, and therefore they have a record of everything that the people in the panel buy.

We did this custom survey asking about

people's migration history on top of that. We combined this with some other data sources, demographics from Nielson of these people, some historical data I'll talk about later, and our sort of -- in fact, we used data on ad intensity of these categories, and we have the kind of subjective coding of social visibility.

So we don't know people's entire migration history. We know where they were born and how long they lived there, where they live now and how long they've lived in their current place, and so we're basically going to drop people who have a large gap between those two windows, so a large fraction of people either have always lived in the same place -- place here means state -- or have only lived in two states.

So basically you want to think of the sample as those people plus if there's kind of six months in between, you think of that as basically measurement error and include them as well, so we have about 38,000 households, of which 10,000 live in a different state from the state where they were born.

I want to carefully tell you what the measure is we're going to look at here because I'm going to show you then lots of things in terms of this measure. I'm going to focus in each category (so we have 230 or so categories that are things like soda or baking powder or whatever) on the top two brands by purchases.

I think of qi1 is the number of times consumer I makes purchases of Brand 1, qi2 is the number of times they make purchases of Brand 2. You can weight this by dollars or various other things, but we're just using number of purchases, and let yi1 be the share of those top two brand purchases that are of brand 1.

I'll call top brands share of the top two brands "purchase share." Then I'm going to define, and this is kind of the key measure, what we call relative shares, so let mu be the average purchase share among non migrants who live in Washington, D.C., and now think about somebody who moved from state R to state R prime. We're going to define that a consumer's relative share to be their purchase share minus the average purchase share they were born relative to the average purchase share in their current state minus the average purchase share where they were born.

This is equal to zero if I look just like someone in the state where I was born. This is equal to 1 if I look just like someone in my current state. If it's equal to .5, that means I'm halfway in between, so this is a measure of if I look at a migrant, how do they look relative to their current and birth state.

So this is a nice summary, and it captures various hypotheses you may have. In a world where there is no persistent effects of where I lived in the past, all that matters for my purchases are the things that would usually be in our models like prices and advertising, availability and so forth. These data should all be 1.

In a world where there's complete persistence, that all that matters for what kind of mayonnaise I like is what kind of mayonnaise my mother used on my sandwich when I was a child. These beta should all be zero, and in the kind of model that we're going to describe, I'm going to write down later they should depend on the number of years I've lived in this place relative to the other place and the age I was when I moved.

So if I was 60 years old when I moved, I had accumulated a lot of capital so I wouldn't converge so quickly. If I just moved to this place, I'll probably still look a lot like people where I was born.

So this is -- I tend to shy away from 3-D graphs, but this is a case where this is actually somewhat useful. These are the betas. Raw data plotted against the age at which -- this is for migrants -- the age at which they moved and the number of years they've been living in their current state.

So clearly these betas are not all equal to zero. Migrants don't look like just places where they came from. They also currently are not equal to 1. They don't look like the people where they live. They are decreasing in the age at which I moved as we would expect and they're increasing in the years since I moved, which we would expect.

So this is just projecting that on the years axis, right, so somebody who just moved from California to Washington is about 60 percent of the way to looking like somebody from Washington so there's this immediate discrete jump. This remaining gap closes steadily but very slowly, so even out 50, 60 years, there's still detectible differences. This is the slice against age, so somebody who moved when they were very young looks pretty close. Somebody who moved when they were older looks farther away.

So all of that is from a cross-section of our data, let's take migrants who moved at different points in time who have lived there different amounts of times for different ages. There's also a little bit of panel component to this data, which is quite short. We can see there's about 220 households in our sample who moved during the two years that we can follow them, and we can watch their purchases before and after their move.

So if the inference is from this cross-section at the right, we should expect a jump of about 60 percent when they move, right, and this is going to be important and I will talk it for distinguishing things, like is this just selection of who these people are, is this really -- is the inference that the cross-section can tell us about the panel really valid.

So if you look directly at the panel, something important to say is we don't know exactly when these people move, so we know people who have moved in the last 12 months and the people that moved between 12 and 14 months, so think about what you would predict.

So these are people who moved between 12 and 14 months ago by month. Their relative share is averaged over these categories, so if our hypothesis -- if the cross-section is right and they jump .6 and if, as you would expect, when they moved within that year were uniformly distributed, so if I could follow the individual, their consumption would jump, but I know you moved in this window but I don't know when. If it was uniformly distributed, this would be linear from zero to .6 and then flat or slightly increasing from .6 on.

So the precision here isn't incredible, but this looks very consistent, and there's no evidence that kind of back here, these people looked all that different from other people in the state where they were born. These are people who moved in the last year. Again, the precision is not incredible, but you see exactly what you would expect to see, flat. Before they moved, they looked just like the people where they were born. After they moved they jump ending right about .6.

So to summarize this kind of descriptive evidence, you see a jump to .6. The remaining gap closes very slowly. It takes 20 years to reach .8 from .6. Even after 50 years there's a significant difference.

You see I don't really show you this in the figures, but if I look at older migrants, the jump is the same when they move, but the gap closes more slowly. This is for, what it's worth, what our model also will predict. Importantly, if you look at the panel evidence, to the extent you can tell, migrants look similar to non migrants before they move, which argues against the idea that we're just selecting migrants who are kind of intermediate in their preferences and that's driving everything.

So you want to interpret this then through the lens of a kind of a habit formation capital stock type model. Everything I've showed you is just data. Once we start putting it in a model, we care about causality and actually interpreting these things. And in particular, there are two assumptions, two kind of key identifying assumptions that you need to buy in order for the estimates that I showed you to be valid.

The first is I want to estimate the impact of you living in California 40 years ago. In order to do that, I need to know something about what was happening in California 40 years ago. I can observe that today, Coke's average purchase share in California is .6, but I don't know what it was 40 years ago. If some place where Coke has .6 today had .55 in the past, we would be understating the extent of persistence. If some places where Coke has .6 today had .8 in the past, we would tend to be overstating persistence.

So econometrically we need expectations that the past equals the present. We actually went and gathered more fragmented data on the past, where we can actually look at this directly, so these are across state categories, purchase share in 1948 to 1968, past purchase share today. There's lots of noise in these data, that this Y axis of these data comes from a little survey that newspapers were doing.

But the coefficient in this regression somewhat, partly we're just lucky here, is almost exactly 1, so the assumption that on average the past share is equal to current shares is confirmed by this data.

Second, we need any unobservable preferences that people have for these brands to be uncorrelated with their migration status. It can't be that they're systematic things about people in California that make them like this brand, and migrants are somewhere in between.

We have three bits of evidence on that. One, I already showed you this match between the panel and the cross-section, so not only did the guys look similar to the people in their birth state pre moved, but it's uncorrelated with the age at which they moved.

Three, we looked at some brands that only were introduced late, so they're saying if there was no frozen pizza in the world before 1980, under our identifying assumption where you lived before 1980 shouldn't matter. Under the selection story, it should because it has information about which type you are which -- and we don't find any evidence. We find where you moved after these things were launched has no effect to the extent we can tell, given the power of this test.

The model is going to have basically two things. One, what we call base line demand which is everything that would go into demand for consumer who had no prior. So, that incorporates all the supply side stuff like prices and availability as well as consumer characteristics that affect these preferences.

The other thing is this capital stock which is just a discounted weighted average of your past purchase shares, right, so we're modeling -- this is about consumption per se, not about what ads you saw, and then we assume that demand was just a weighted average of these things. So the parameters we're going to estimate are the weights on current stuff relative to capital stock and the formation of capital.

So these estimates are almost just literally mechanically repeating that figure. We designed this model such that it's a very simple model that basically captures exactly that slope that you saw with respect to years, adding controls, adding this kind of discounting process to separate these two things, but you should have expected this almost has to be .6 given that's that jump that you saw.

The discount factor we estimate is .975, so that means for example if I get a unit of capital today, it takes 27 years for half of that unit of capital to decay. These are the fitted values from that model, so you can see it's kind of fitting well, the rough pattern that we saw, and you can look at the residuals and they're not systematic.

So then we do some other things that I don't have time to show you in detail, but we show that the brand capital is more important in categories where advertising is high, in the sense that that weight on the brand capital term in your utility is bigger.

Basically that means the jump when you move is smaller for these categories. Brand capital is more important where social visibility is high. This implies we think about these kind of counterfactuals and big first mover advantage, so for example if one brand has a 10 year head start, they would need to discount price by 40 percent for 25 years to catch up to their competitor.

These things last for a long time, and we think this can rationalize something that was observed in the other paper, which is who entered Detroit first in the ground coffee category in 1900 still seems to have a big effect on who is ahead today, despite wars and depressions and recessions and advertising campaigns and all kinds of stuff intervening.

So I'll stop there. Thank you.

DR. NEVO: Thank you. Our discussant is Matt Weinberg.

DR. WEINBERG: Thanks. I'm not going to spend a huge amount of time summarizing. You guys just saw it, so this paper picks up with an earlier paper by a subset of the coauthors on this one in the JPE which finds that a lot of the cross-sectional variation in market shares today are explained by who entered first a long time ago, and that kind of begs the next question of what exactly is causing this persistence and this paper does a nice job of explaining what that is.

It's able to separate out these kind of supply side variables like variability and cost advantages from brand preferences using the information on migrants, and again the big results were that it looks like about 60 percent of the gap between the average purchasing pattern is explained by the supply side variables, and 40 percent is explained by persistent brand preferences, and that gap closes over time, but it closes slowly.

So again several reasons were given for what might be the root of this brand capital. A few things that were mentioned were -- habit formation was actually having consumed the goods in the past, that that causes the persistence, or it can be exposure to advertising or learning from others, although these peer effects, having other people see you consume things might matter as well.

I'm wondering if there might be a little more on the paper that could be done to try to differentiate between these different explanations, so I was wondering if perhaps you could identify some products that are typically purchased when you're old, and I identified three, maybe they're kind of funny products, but typically people haven't given a lot of though or consumed denture cleanser, hair dye and incontinence products when they're young.

So I was wondering if the effect holds for these products. If it does, then potentially it's due to advertising. This isn't perfect because you're probably not really high advertising markets or learning from family or something else that's causing the persistence in brand capital.

So I'm wondering if you get the fact that migrants look like people from which they came for this type of good. I'm not sure if you have the age of the children of the households in the Nielsen data, but if you do, perhaps you could look at people that have kids but have them after they've moved to do something similar.

In summary, I think I've learned a lot from reading this paper, and I appreciate the fact that the authors went out and collected data to answer this question. I'm kind of curious about how the respondents to the additional survey that collected migrant status compared to the rest of the people of the Nielsen Homescan data. Do they look similar or are they older, younger or wealthier? And I found the results very convincing and enjoyed the paper.

Thanks.

DR. NEVO: Let's move on to the next paper, and then we can have questions on all the papers together, so Kate Ho.

DR. HO: I'm going to talk about something completely different, which is medical care. I'm looking for physician responses to financial incentives, and this is joint work with Ariel Pakes.

So there are really two motivations for this paper. The first relates to the U.S. health reforms from earlier this year, which includes among a lot of other things new provisions to give physicians financial incentives to control costs in the Medicare and Medicaid programs.

So unlike in the current system where a physician providing a Medicaid or Medicare, receives a fee for service payment, you provide a service, you get a payment for that service, the reforms will set up organizations called Accountable Care Organizations which will be groups of providers that are eligible to share in any cost savings they make from the Medicaid and Medicare programs.

There are also going to be pilot arrangements under which physicians providing Medicaid services receive bundled payments for episodes including hospitalizations. So, the goal of these kinds of provisions is clearly to give physicians incentives to control costs, hopefully without comprising on the quality of care they provide.

It turns out that similar cost control incentives are used currently by Health Maintenance Organizations, HMOs, for privately insured enrollees in California, so there's an obvious opportunity here for us to try to understand how physicians respond to those kinds of incentives. So, that's what we're doing here.

There's been a previous literature that looks at these kinds of issues. There are a lot of papers that document lower costs in HMOs compared to other types of insurers, but in general they don't look in any detail at the mechanism used to reduce costs. In this paper we're looking at one specific mechanism. We're asking whether patients whose physicians have a financial incentive to control costs receive care at lower priced hospitals than other patients.

The second motivation is much broader than this. It goes outside of the U.S. health reform. There's a big previous literature that uses hospital discharge records to estimate models of hospital choice, and these models are important for all kind of regulatory analysis, for example to predict the price effects of mergers, hospital mergers or to understand hospital incentives to invest in new technologies.

In general, the way these papers work is, first of all estimate a model that asks how much decision makers value each hospital, and then run counterfactuals asking how much that valuation would change after the merger or the investment. But these previous papers in general ignore the impact of the price paid by the insurer to the hospital. So, we're going to address that issue. We're going to ask whether the hospital choices are ever influenced by the price paid by the insurer to the hospital.

So I've got about 15 minutes, I'm going to try to get through this. I'm going to give you a very quick overview of what the market looks like and of the model. First of all, it's important to explain why we think choices should respond to hospital prices, and then I'm going to talk briefly about how we're going to go about estimating this price sensitivity.

I'll tell you something about the data and then in a bit more detail about the model. There are two methods we're using. One is a multinomial logit analysis that is very close to the previous literature on hospital demand, and then we're developing a methodology based on inequalities to address some of the problems with that methodology, and I'll show you some results at the end.

So a little bit about the market we're looking at. So this is the California medical care market in 2003. We're focusing on HMOs, Health Maintenance Organizations, which cover something like 50 percent of the employed population in California. The seven biggest HMOs covered 87 percent of the HMO market, and we're including six of those seven, all of them except Kaiser.

So how do physician contracts work in this market? Well, other than Kaiser, outside of Kaiser, the model that dominates is the California Delegated Model, under which HMOs sign non exclusive contracts with large physician groups. And there are two payment mechanisms for these groups.

The first is a capitation payment system under which the physician group receives a fixed payment per patient to cover the services provided to that patient. It turns out that under these capitation payment arrangements, 89 physician groups have incentives to control hospital costs, and also that these incentives in general are passed down from the physician group level to the individual physician level.

The alternative payment arrangement is a fee for service contract, and these much simpler contracts in some sense don't generate these incentives for physicians to control hospital costs. So, that's useful for us in this analysis. There's capitation payments under which physicians have an incentive to control hospital costs, fee for service contracts under which they don't.

So how are we going to use this analysis? We're using hospital discharge data for California in

2003. We're focusing on women in labor; they're going into hospital to give birth. Unfortunately, the data set does not tell us anything about the patient's physician groups and very little detail on the compensation schemes used to pay these physicians, but we do observe each patient's HMO and the percent of each HMO's payments for primary services that are capitated. There's a lot of dispersion across insurers, from BlueCross/BlueShield at the low end, who banked 38 percent capitated payments, Pacificare at the high end, 97 percent.

So the questions we're going to ask are, first of all: Are hospital choices influenced by price? And secondly: Does price matter more when the patient is enrolled in a high capitation insurer? By assumption, does price matter more when there are incentives to control costs?

So here's a little overview of the model we're using. The idea was to estimate the utility of the combined agent making the hospital choice. That choice is made by a kind of composite agent, the patient, the insurer and the physician. We're estimating the utility of that composite agent. So, where is the utility generated when the patient goes to hospital is going to depend on the price paid by the insurer to the hospital.

The second term there is an interaction. Let me talk about that in a second. The third is distance from the patient's home to the hospital, which has been shown to be important for hospital choice, and there's an arrow term at the end. I want to talk a bit about this interaction term G pi but this is the key for us to get these estimates right and to believe that the price coefficient we're estimating is right.

So that G pi term is an interaction between measures of patient severity of illness and the quality of the hospital. We think it's important to get this interaction term fully flexible for a couple of reasons. First of all, because we think that hospitals are likely to have higher quality for some sickness levels than for others; secondly because the preferences of a decision maker for quality are likely to differ across severities.

If we don't get those things right, we don't control for them fully, we're going to have a biased price coefficient, so it's important to get that right here, and I'll show you in the two methodologies the extent to which we were able to do that, but again the questions we're asking, first of all: Is that price coefficient negative? Secondly: Is it more negative when the insurer capitates physicians?

So the data set: I said we have hospital discharge data from California in 2003. It's a census of hospital discharges, so we have an observation for every discharge. We're looking at privately insured HMO enrollees, women going into hospital to give birth. At the patient level, we observed the name of the HMO, the name of the hospital, a lot of detail on diagnoses and procedures and something about the price paid, although it's not a perfect measure of price. I'll tell you about that in a moment.

At the hospital level we observed the average discount. That's also going to be an input into the price paid, the hospital's location, teaching status and detailed information about the services provided by the hospital and financial statements.

The price variable is going to be an important input into this analysis for obvious reasons. Unfortunately, we don't observe the exact price paid to the hospital. Instead we observe a list price, which I think is equivalent to being a hotel rack rate. Every year the hospital publishes a schedule of its list prices.

Very few patients actually pay those prices. Instead what happens is each insurer sits down every year or two with every hospital and negotiates a discount from those list prices. We observed the average discount at the hospital level.

So we're going to calculate a price measure in two steps. First we calculate an expected list price, which is an average list price for ex ante similar patients at the relevant hospital, and then we're going to assume, at least for now, that the discount is fixed across insurers. There's a lot of things we want to do in the future to make this estimator better, and one of the things we will do is to change that assumption, but for now we're defining the price as the expected list price multiplied by one minus the average discount.

Here's something about the discharge data we're using. We have 88,000 patients, 195 hospitals. Teaching hospitals are a big deal of course. 27 percent of discharges are from teaching hospitals. The average list price, once we've interacted with a discount, is about $4,000. And I'm showing you also average length of stay and some probabilities of adverse outcomes, which is low probability of events, since very few women in the U.S. die in child birth. The probability of transfer to an acute care setting is about .3 percent on average. The average probability of transfer to a special nursing facility after giving birth is just one and a half percent.

I want to give you an idea of how these prices and outcome measures differ by patient type, to give you some idea of the variation in the data. So first of all, by age. Not surprisingly, most of the women giving birth are age under 40, and you can see that the average price for their procedures is significantly lower than that for older women. The probabilities of these adverse outcomes are also lower for younger women.

The second panel is showing you variation by patient severity measured by something called a Charlson score, which is a clinical index developed by physicians that assigns weights to comorbidities. A higher number means a sicker patient. The vast majority of women in our data have a Charlson score of zero. So, they're really not sick. They're just in hospital to give birth.

About 1,800 have a Charlson score of 1. About 80 have a Charlson score greater than 1. And you can see that again with variation in the data is very much intuitive so price increases significantly as we moved from group to group. The probabilities of these adverse outcomes also increased significantly from group to group. So that's just to give you an idea of what the data set looks like.

Let me move on before I run out of time and tell you about the methods we're using. So first we do a very standard demand analysis. This is very similar to the previous literature. We're estimating by maximum likelihood. The equation for estimation I've written again at the top of the slide. It's essentially the same as you saw before.

The big issue for the logits is that this interaction term that I said was so important to control for price endogeneity is defined much less flexibly than we would like, and there's essentially a feasibility issue here. The approach we're taking to try to control for price endogeneity is to put in as much detail in that G pi term as we can. And in the inequalities, you will see that essentially means having something like 16,000 interaction terms in this equation, and it's just not feasible in the logits to do that.

So we're defining it in much less detail. That means there's a caveat going into these results that I'm going to show you in a minute, that is, we're expecting the price coefficient to be biased upwards, and that's what we see.

So when we estimate the logits for all patients, we get a positive significant price coefficient. That's the panel on the left. The next thing we do is to split the sample. We look at half the sample that's relatively less sick. When we do that, we get a negative marginally significant price coefficient consistent with the idea that this unobserved quality that we're not doing a good job of controlling for matches more for sicker patients.

So we get a positive significant price coefficient for sicker patients, negative and marginally significant for less sick patients. The distance coefficient, as I said, we know distance matters. It clearly matters here as well.

The next thing we do is to allow that price coefficient for the less sick patients to differ across insurers, and this is the first step we're taking essentially to test our ideas about whether capitation payments matter. So you can see that I'm allowing the price coefficient to differ across insurers here.

Here's a list of the six insurers I'm looking at, and I've ranked them in decrease in percent capitation from Pacificare at the top to BlueCross at the bottom. The estimates are on the far right of the slide. You can see the price coefficients are negative for the top four insurers, significantly negative for two of them, and positive for the insurers at the bottom.

So this at least is consistent with the idea that price matters, but only when insurers are giving physicians an incentive to control costs. In terms of magnitudes, the distance elasticity implied by these estimates is minus 2.7. The price elasticity for Pacificare, which has the most negative price coefficient, is minus .25. Price matters significantly, but the magnitude of the effect here is very small. Again this is consistent with the possibility that there are really price endogeneity problems here.

So the next thing we're doing is to put together a method that is going to deal with that price endogeneity. And I want to take a couple minutes to explain how this works. We're writing down an econometrician prediction of the utility generated when the patient goes to a particular hospital, so you're seeing essentially that equation at the top as from before.

There's a price term, a G pi, this interaction term that we need to control for and a distance term, but if we change, we're now defining patient severity at a more detailed level than we were able to in the logits, and the G pi is going to be a fully flexible interaction between dummies for these patient severities and hospital fixed effect. So, we're defining everything in much more detail here.

We end up with 157 hospitals, 106 severity groups. That's where the 16,000 fixed effects or 16,000 coefficients comes from. Given that we've defined it in much more detail, we're making an assumption which is that that G pi terms absorbs all the endogeneity problem. It absorbs all unobservables known to the decision maker that effect hospital choice.

Then the remaining unobservable is such that the expectation of that error conditional on instruments is zero, and that gives you a utility equation at the bottom of that slide, which is the utility observed by the decision maker and that's used in the choice.

Then how does that methodology work? Well, we're making a simple assumption which is that for every patient who goes to hospital H, the utility generated from the chosen hospital is greater than or equal to the utility generated had she gone somewhere else, and that's what the inequality at the top of the slide says.

A little bit of notation, this W ih, H, H prime is the difference between the utility when that patient went to hospital H and the utility had she gone to H prime. By assumption that's non zero. And the intuition for what we're going to do here, we're going to find all pairs of patients who went to the same insurer and had the same severity but went to different hospitals. Patient ih went to H and could have gone to H prime. Patient ih prime visited H prime, but could have gone to H.

We're going to write down the inequalities and add them together. Those G pi terms, those complicated interaction terms are going to difference out because we've chosen those patients carefully such that they have the same severity and the same insurer, and that's the key to this methodology. Once we've differenced out those G pi terms, we don't have to estimate them anymore. We can define them in much more detail than was possible in the logit analysis.

Then we're going to take expectations on a data generating process to address the measurement term. I'm not going to talk in detail about that. I'm going to show you some results from that inequality analysis. Again here are the six insurers we're looking at. At the top is the high capitation insurer, Pacificare; at the bottom, the lowest capitation insurer is BlueCross/BlueShield.

First I'm showing you the results using a very small set of instruments. You can see that we're getting a range of values for the price coefficient that are consistent with the inequalities we've generated. Quite often just using these instruments, we have either an upper bound or a lower bound for the range but not both.

Still there's a consistent story here, which is that for the top three, it shows that for the high capitation insurers, the price coefficient is clearly negative. The upper bound for that range of values is negative. For the bottom three insurers, that's not the case. We can add more instruments and you can see that now we're getting a well defined lower and upper bound for the price coefficient for every insurer, and the same story holds.

So for the insurer that makes a high proportion of capitated payments to physicians, that gives physicians an incentive to control costs, the price coefficient is negative. Price matters in a negative way. For other insurers, that's not the case.

Very briefly in terms of magnitudes, I showed you for logits that Pacificare had a price elasticity of a demand of about minus .25. For the inequalities, the magnitudes are much larger. The elasticity now is minus 4.1. Health Net, which has a lower proportion of capitated payments, the numbers are smaller, but still the magnitude for the inequalities is large, with an elasticity of minus 1.9.

There is some comparison to the previous literature. Gaynor and Vogt have a paper putting a price index into the utility equation, and they get an average price elasticity of 4.9. So, even bigger than what we get.

That doesn't necessarily mean that I think these magnitudes are realistic, particularly the 4.1. This is something we're working on. But in general the message is clear that the price seems to matter. It matters particularly for high capitation insurers, and when we have a method that deals with the price endogeneity that we're worried about, the magnitudes of the effects are large.

So quickly to conclude, what we're trying do here to estimate the preferences of the agent that determines hospital choice and identify whether physician incentives affect price sensitivity. We've used two methodologies. Both of them are showing that price does affect hospital choice, and that it matters more when the insurer is capitating a large proportion of physicians.

We have these inequalities method, which we think is kind of cool, it allows us to do a lot of things. In particular we're addressing the endogeneity concerns much more fully than was possible with the previous methods. There are a couple other advantages that I don't have time to talk about.

There's a lot more that we're trying to do here in terms of developing this analysis, and if anybody is interested, I can talk about it for a long time later, but we think that the results we have so far have real implications for the impact of the U.S. health reforms on costs and for regulatory analysis more generally.

So thanks.

DR. NEVO: Thank you, Kate. The discussant is Keith Brand from the FTC.

DR. BRAND: I will try to make this as brief as possible to keep us on time, so to summarize what we just heard very briefly, this paper examines the price sensitivity of a composite insurer, physician, patient age on their choice of hospital using two approaches, a conditional logit, and a inequality analysis. I think Kate's described the motivation for the inequality analysis, so I don't want to go over that again.

The results are intuitive in that there's basically increasing price sensitivity in the degree of capitation of the insurer in both. In the logit model, this holds only for the sicker patients, but not the least sick patients.

And in the inequality methods, it holds on average across the population, and I'll note in the paper that there are small differences in the results so far when you divide the simple into less or more sick patients. So I think this is very important stuff for the policy questions that Kate has already outlined, and for me in particular because I'm dealing with these issues quite frequently in my work at the Commission.

So I don't have many comments. The biggest comment is kind of coming from the perspective of someone who has an interest in this approach and would like smaller information on the paper on how to make the -- on how to assess how you define the groups so that you know you've got the group definition correct.

So but one quick point, in the logit results we saw the variation in the results between the less sick and the more sick patients, and they pose the question in the papers: Is this a result of a systematic variation in preferences or is this evidence of endogeneity bias? And the rationale in the paper was to put more weight on the latter, that basically the insurer pays the price and the insurer's willingness to pay for a fixed utility benefit shouldn't vary across patients.

I guess my initial response to that is this is a composite of the preferences of the physician and the patient also should matter. Even though the utility benefit may be fixed across patients in the model, it's probably not fixed in the underlying data generating process. There's still the market utility from an increment of quality. It is probably going to be higher for the more sick patients than the less sick patients, and I guess at the end of the day, it wouldn't surprise me that you would see some systematic variation in preferences, even in your inequality analysis.

So the inequality: my big comment is on how one defines the severity in price groups here. There are a number of trade-offs that Katherine and Ariel have outlined in the paper. You define the severity groups in a way so they're refined enough so that you can plausibly say that any remaining price variation is going to affect choices, but there's no correlated variation and unobservable quality that may affect choices as well, which would bias your price estimates.

So you want to set up and refine the groupings, the severity groupings in a very refined way, but not so refined that you wipe out all the price variation so there's some trade-off on that.

Within severity groups, you define price groups so then again you can plausibly say all the variation in price within severity groups affects choice, but is not correlated with unobserved quality, and there are a number of --there are some measurements issues here.

If you define the price groups too broadly, there are aggregation issues. If it's too narrow, there are some measurement error concerns. And again you want to pick the price grouping so that you maintain some variation within the severity group. That's a poorly worded bullet. You want to maintain price variation within severity groups, not price groups because you're going to compute an average price within price groups.

So again my question in going over this paper was: How do you know that you got the right balance? How do you know that you're not wiping out so much of the price variations that you're getting kind of a precision on the bound? How do you know that you've defined the severity groups in such a way so that you really can say that there's no endogenous variation?

So in the paper they talk about -- so the key distinction kind of, this is maybe too crude a description, so basically within severity groups, they look at the variation of co-morbidities across price groups, and the question is: Does the co-morbidity or does the variation in co-morbidity within price groups explain choices. And they relied on the opinion of outside experts which basically said no, which is of course very valuable.

They did an analysis of variance on the outcomes, the probability of mortality, the probability of transfer to another facility. Does moving from severity groups to price groups explain additional variation in these outcomes? And if it doesn't, then it's probably reasonable to say that you've wiped out all the endogenous variation that's coming out of prices.

So I guess what I would appreciate more is how much -- what is the incremental explanation, what is the incremental -- how much of the variance on outcomes is explained by prices? And you say that it's small, but I guess I would like to know how small is small enough for you to be comfortable with that.

In addition, once you've defined these groups, they look at the -- how much of the overall price variation is explained additionally within price groups going from severity groups to price groups, and it's 12 percent here. So, that doesn't strike me as implausibly small, but I guess at what point would you have considered that too small to say, "Yeah, I have the right amount of price?" Or if it's too small, would you back off or further refine the pricing groups so that you can see that you're generating more price variation?

Alternatively it seems plausible to me that if instead of 12 percent, you have 60 percent, you might be concerned that maybe the definition of the severity groups wasn't sufficiently narrow to actually filter all the price endogeneity. So, there are a lot of moving parts it seems to me at this point in setting up these things, and it seems very critical because my guess is the results are very sensitive to exactly how you define this stuff.

So that's my big broad comment. Let me make three smaller points, if I can get this in. So you have this constant variance on the bounds, and I'm putting variance in quotes because obviously we're not talking about classical inference here with a covariance matrix, but you know get some commentary results on the CIGNA observation. You get a positive lower bound on the price coefficient, and you know that, well, this could be attributable to the fact that we have a smaller sample size for CIGNA and so you can get a noisy measure.

So then it occurred to me that if that's a relevant thing, then maybe we should be thinking of some concept of dispersion around these bounds. And I'm not sure exactly how you do that, but the most obvious thing that came to my mind was bootstrapping, presenting a bootstrap distribution on the bounds.

Second, price elasticity comparison to logit, because in the paper you dump the lower bound on the two pairs, what they are into the logit model, and I wasn't sure why you would want to do that.

I'm guessing you took this into account but it wasn't explicitly in the paper, that you normalize the price coefficient differently so in the inequalities model, you normalize it by the distance coefficient, and the logit obviously is normalized by pi root 6.

But what occurred to me is that you could simply stick the price change into the equation 14 and count the number of switches, and that would be something a little more straightforward.

Finally, as Katherine noted they see hospital prices at the hospital level and not the hospital insurer level and pose the question at the end: Could our results be explained by simply this measurement or could it also be explained if higher capitation insurers negotiate smaller discounts, although the sequels only present some regression results that suggest otherwise?

I think my standard intuition is that it's got to be true, all else equal, that hospitals should be compensated for bearing additional risk. So, how could one would address that in the regression model? The most obvious thing that comes to my mind is that the variation of the relative bargaining positions that may explain why some plans are higher capitated and also able to pay hospitals at a lower rate.

So I'm not sure how exactly one could address that, but I guess you would have to look at the regional variation to look at how much competition each hospital has or how many close ups each hospital has and maybe account for that in some way.

So those are my comments.

DR. HO: Do you want me to comment now?

DR. NEVO: If you want like in 30 seconds, if you think it's really important to do it now and not later. Maybe you will get a chance later.

The last paper is Nathan Miller.

DR. MILLER: I want to talk today about modeling and estimating models of spatial competition. This is joint work with my coauthor Matt Osborne, and I should say at the outset that my views that I'm expressing today are my own and shouldn't be purported to reflect those of the Justice Department or the BEA.

I'm going to start with I think the uncontroversial statement that firms in many industries are geographically differentiated. I've thrown a couple up on the slides, and this fact that there has been some competition in these settings has motivated a number of seminal theoretical papers.

I think folks are familiar with the Hotelling line. This stuff goes way back. Though there hasn't been a lot of structural work on the industry, and what Matt and I wanted to do is construct an estimator that lets us estimate the underlying parameters of demand and supply in these models. And do so using variation that as economists, we may reasonably be able to, using data that we can get our hands on.

So, for instance, the estimator we're going to introduce can estimate the parameters of fairly rich models using data on maybe regional average prices or total consumption and production in various areas or firm level data. We think that the structural estimation of these models is interesting and could be used in a number of different settings.

First of all, we could get a grasp simply on how much firms are indeed spatially different and how much local market power exists in industries. And also it could enable new counterfactual policy experiments. One could use the model to conduct hypothetical monopolist tests and construct geographic antitrust markets.

An interesting question might be how carbon taxes or gasoline taxes might affect localized market power on these industries or how imports and tariffs are likely to affect consumers across a large nation or you can plug this estimate more into a dynamic model and start asking questions related to entry deterrence or related topics.

So I want to start by motivating the paper with a simple question: Why is this challenging? The most obvious way to estimate the cost of transportation is simply to observe the distribution of market shares and how they change over space. So, on the slides I put up a little plot. The star is meant to represent a firm, and one could imagine the shades of blue representing the market shares of the firm.

This firm captures high market shares among nearby consumers, and the market shares attenuating in the distance. Maybe this is a movie theater or a gas station or something like this. If you have these data, then one can write down a model and simply select a parameter calling it transportation costs and rationalize the distribution of market shares.

Matt and I like to talk about a data availability problem, which is that distribution of market shares are typically not observed in the data, at least the data that as economists we have access to, and we're not actually aware of any studies that make use of distributions of market shares.

More commonly, it would be firm level shares or prices. I also want to note that in some industries, especially business to business industries, firms exercise spatial price discrimination. For example, they might charge higher prices to nearby captive demand. When spatial price discrimination is used, then one would also need to be able to account for the spatial distribution of prices, which just makes this problem more pronounced.

So the paper really has two parts. The first part is we write down an estimator for a model of spatial price differentiation, encompassing spatial price discrimination that can make use of different data, different balances of observation. And we think that the estimator potentially allows us to -- allows econometricians to extend their models, to extend the estimation of models to settings that previously would have been too demanding or too hard to do. We're going to get some conditions whereby the estimator is consistent and not systematically normal.

The second part of the paper is an empirical application to Portland Cement. Cement fits the model well in some sense. It's relatively homogenous aside from the geographic component. I'll talk more about that later on, but our main goal is to show the estimator works well in this one real world example.

We provided fits that are I think pretty impressive both in sample and out of sample, and it also lets us highlight some of the counterfactuals one might be able to do. So, it allows us to do merger simulation in which we show how merger harm is distributed across California and Arizona and how different divestitures affect not only the total harm, but the geographical distribution of harm.

The main methodological insight here is that you can essentially use numerical approximations to equilibrium to relax the data requirements of the estimator. So basically leveraging the information in some supply and demand model for a given parameter vector, a parameter vector, i.e. compute the prices and the market shares that characterize equilibrium given the model, given the candidate parameter vector.

With the disaggregate shares and prices one can aggregate equilibrium predictions at the level of the data. So, for example if one has data on average prices in California, after computing equilibrium, one might have average prices for each consumer in California, and you just average that up to get the -- to construct the aggregate prediction of the level of the data.

Since this is repeatable, we can do that for any parameter vector, so we can select the parameters that match the predictions to the data. Intuitively the way this estimator works is that you have some sort of nested logit in which you minimize an objective function, you have to compute equilibrium and then aggregate the equilibrium predictions to the level of the data.

The key assumption for identification is that when evaluated at the underlying population parameter vector, the differences between the predictions of the model and the underlying data are due to measurement error. For example in an application on data that are collected by the U.S. Geological Survey, sometimes plants don't report their information to the USGS, so that creates a measurement error.

We're attributing differences between the predictions and the data to measurement error, and we're assuming that that measurement error is going to be orthogonal to plant locations and the cost and demand shifters. Given those assumptions, you can derive what essentially amounts to a multiple equation nonlinear least squares estimator. The left-hand side of this are essentially the data. The twist is the right-hand side of the model predictions are computed, so the right-hand side of this is based on equilibrium computations rather than data itself.

Intuitively each equation in the least squares estimator matches a times-series of data, for example, the average prices in California to the corresponding prediction, equilibrium prediction.

I'm going to start by describing a little bit an economic model, and then we'll talk about how to take this to the data. We start with a notion of geographic space which is some area, and plants have fixed locations. Consumers exist over the space, and then we're going to introduce the notion of a consumer area, which is a subset of the space. And each firm is going to be able to set a different lower price to each area, and consumers are going to bear the cost of transportation.

The consumer areas allow us to build in spatial price discrimination. So, each specification determines how much of the area us discriminated. For example, if there's only one consumer area, then there's no discrimination, but if there's lots of areas, you have discrimination.

This is an example of one geographic space. There are three consumers areas and two firms. Each firm is setting three different prices, a different price to each consumer area.

So the supply model is fairly straightforward. Firms are maximizing variable profits, which is just the price times quantity in each consumer area indexed by N, and then less variable costs which is just you can integrate up over a plant specific cost curve. We need the cost curve to be continuous differentiable, but potentially you can capture things like increasing marginal costs or capacity constraints flexibly.

Demand within a consumer area we're going to model using a conventional discrete choice system. The indirect utility is just going to be a function of the price that's charged in the area, the average -- the average distance between the consumer area and the plant.

Now, if the error term is logit or nested logit, then one gets analytical expressions for market shares, and that's going to facilitate the computation of equilibrium quickly.

You get standard first order conditions here, and this is really the key to the model, and one can characterize equilibrium as just a mapping from the parameters of the model into a vector of prices such that the first order conditions hold. We're going to assume that this equilibrium is unique and that it exists. And I'll come back to this: Why it was so important?

So given the structure of this model, what we want to do is we want to recover the underlying parameters. I'm going to denote the endogenous data as a vector YT. This includes average firm prices or production or anything else like that.

I'm going to denote the aggregated

equilibrium predictions as Y tilda. That will be a function of the parameter vector as well as the vector X, which will include plant locations and cost and demand shifters and things like that.

The estimator takes the following form: I'm just going to minimize the squared deviations between each of the equations and the equilibrium predictions, potentially weighting where appropriate between the different equations.

As I said before, this really amounts to non-linear least squares, and once you get here, this is textbook. You can open up Green, and there's a description of this. The twist is the right-hand side, the Y tildas are computed, and what I do is we just basically select a price vector that makes the first order equations almost whole, at least very precisely to a small -- so errors are small, so we use a tolerance one minus 13. To actually define this non linear, we use DFSANE, and we can end up doing an iteration in two to ten seconds or so.

If a unique Bertrand Nash equilibrium holds, exists, and the population parameters vector is identified, then the NLLS estimator is asymptotically normal. A couple comments: One is that there is the uniqueness of existence so the logit is in multiple firms but it's not generally a property of the models we're looking at.

The second is that the population parameter, the aggregation procedure can obscure identification. So, even if one had disaggregate data, potentially you use the identification of the ideation process. Matt and I talk about how you can potentially test some of these assumptions empirically.

The empirical application is Portland Cement. Cement is defined as a finely ground powder. You put it with water. You get ready mix concrete. You ship it by truck. Consumers pay a transportation cost. Contracts are individually negotiated.

This is a map of the area we looked at. Plants are in blue. Imports flow in through San Fran, LA, San Diego and Nogales. This model also captures the foreign imports. There's very little inflow or outflow from this area to other domestic areas, and so you really you do get a geographic space in the sense of the model.

We use a marginal cost curve that bends upwards at some point that we estimate. Demand is specified with nested logit where we put the outside goods in a different nest. We use 90 counties within Arizona, Nevada, California to specify price consumer areas, which of course has a fairly fine price discrimination, and we model importers as being a competitor fringe.

The data we end up using are average prices in three regions: Total production in those regions, total consumption in each of those four states, and we also make use of the little information on the cross region shipments. So, we end up with essentially ten non linear equations over 21 time periods spanning 1983 to 2003.

Here's the model. We can see in panel A we look at the regional consumption. On the left-hand side is the data. On the right-hand side is the model prediction that the estimates have produced, and we can see we've explained about 93 percent of the variation.

Panel B looks at production. We're explaining 94 percent of the variation there. Panel C is the 82 percent of prices, and panel D is this out of sample, which has 98 percent of the region, cross region shipments.

We're able to do some neat things like this. This is the distance the cement is shipped over across the space. This is plots of the business so there's a plant here with a star, and you can see how its prices go down to consumers that are more distant, and as do its market shares. This plant, which is just north of Phoenix, seems to have a fair amount of low price market power.

Here you can see that that merger harm is concentrated around LA and Phoenix diminishes elsewhere. And in Map B, we've examined the merger harm under one potential divestiture plan in which we divest one plant in the LA area, and one can see that it mostly mitigates harm in southern California, but not much in Phoenix.

Let me go through this just to finish. One is that our mix uses the estimator in a static model, but the estimator could also define stage game pay offs in more dynamic routines, so essentially one could plug it into maybe an estimator by Bajari, Benkard, and Leven and things like that.

I think some of the interesting questions that would enable economists to answer is to look at firm location choice and ask the question how firms should ultimately locate to deter entry or whatever. You would have to solve the state-space problem to get that done, but we think it's potentially an interesting extension.

Second, there's a parallel here to estimators for product space differentiation. And you can intuitively take BLP: When you run BLP, generally what you get is you fully observe prices versus market shares at that level, but you don't observe all characteristics, for example quality.

In our model, you fully observe the characteristics, but in the data you don't observe all the prices or market shares. Instead what you observe is aggregated. In both cases we're using numerical techniques to cover unobserved metrics, and then enable us to make sure there's a minimization of an objective function.

So that's it.

DR. NEVO: Thank you. Our discussant is Mr. Cement himself, Allen Collard-Wexler.

DR. COLLARD-WEXLER: I want to start off with just some review of why we think spatial markets might be important and difficult to analyze.

So the central issue that we have with these spatial markets is that essentially there is market segmentation, but the markets overlap with each other. So, even though I might be a cement plant over here, and I might only compete with cement plants that are around me, those plants compete with other plants that are located further away and so on.

So rapidly, kind of solving out equilibrium in spatial markets has this problem of dealing with neighbors' neighbors, and the state space becomes huge. This is an unfortunate problem, given a lot of markets we hear about have some spatial segmentation. This is clearly true for cement, and it's true for a lot of other bulk commodities, things like coal or chemicals, for instance, electricity as well in a much more complex way. These are important market power issues to get a handle on and they're difficult.

What I've been trying to summarize on the paper and what I want to focus in on is what specifically this kind of paper adds. It has some nice features.

One of the features is this is a model of spatial price discrimination, so there isn't just a price that the cement plant charges and then people pay transportation costs to the final location. There's a price that the cement plants charge to everybody in a specific county, so instead of having to deal with 14 prices, there are 14 cement plants here, they have to deal with 14 times 90.

So they have this huge increase in the number of prices that they have to account for in the first order conditions, and there's some fairly large stuff that they have to actually do to actually get this to work, and it's fairly impressive. And I think the reason why it's so much more complicated is so they can deal with spatial price discrimination. So, it would be nice to have more evidence on how spatial price discrimination changes the predictions of the model and how it affects market outcomes.

In particular, there's a lot of spatial price discrimination in a lot of markets. Some of the times, you don't realize it because everybody gets charged the same price, but there may be different transportation costs, so that's implicitly spatial price discrimination. We just don't see it that way. So I think it would be nice to emphasize this particular feature of the model.

The second feature is for demand and cost estimates, the model is actually getting very reasonable answers, so on the cost side, they find fairly reasonable estimates of transportation costs, which is a nice check that the estimates are doing a good job. Also, in terms of getting aggregate demand elasticity right, they find demand elasticity of .16 percent and at the firm level, 5 percent or so.

So, we think that cement markets are very inelastic demands at an aggregate level, and they're finding this, and this is to contrast with other work which finds elastic aggregate demand for cement, which just seems wrong. I think this is a nice way of testing that the model is giving reasonable predictions.

So going on to the one conceptual thing going on in this model that's maybe controversial but also usual. So, the typical approach in IO is, "Well, let's look at the actual shipments or let's get very micro data to analyze the problem of transportation costs." Often that data is really hard to get, and on top of that, often the moments that you're going to use for estimation, like the average distance traveled might be measured incorrectly.

So maybe I'm using the distance just in miles, but sometimes I'm transporting along the highway. Sometimes I'm not. So, it might be mis-measuring transportation distance. So, instead of focusing on micro moments, they're using these aggregate moments, and then consumption in different areas are right. Sometimes we're paying attention to those aggregate moments rather than micro moments, and it might tack down the estimates and give you more reasonable results. So I think there's some value here in being able to do that. This is often the only data that's readily accessible.

This method could be used for other cases where we're trying to get the kind of topology of trade costs. For instance, for cement, there's a big difference between water costs and land costs, and for other markets, that kind of difference in transportation costs might be something we are wanton for.

I think they have a very rich model and I think they're looking for what the most useful applications of the model are. There are two things that are hard with looking at dynamics for cement. One is we don't see a lot of entry and exit and the second thing is the state's base is huge, so it's the configuration of all the plants in the entire market. And that becomes difficult. So where useful, it's hard to know where to take it.

I think instead what this model does better than any other model I've seen is get at spatial price discrimination, and so can we say anything about are there welfare effects of allowing versus not allowing spatial price discrimination? What are the overall effects in this market? And essentially you can do this very easily, and it's not something that empirically you have a lot of evidence on, and for the price discrimination in general, the welfare effects are typically ambiguous of having it or not having it.

So it's something we fear is not going to give us a clean answer, and then there's some literature on different forms of spatial price discrimination like basis points on steal that maybe you could also refer to.

Then the other issues are also a large role of international competition in this market, but the topology of competition for cement is very weird because it hits coastal markets in a very different way than interior markets.

So this model might also have something to say about it. Other people haven't worked on this before. So, I think that would be an interesting way of tying in those trade costs.

Thank you. That's it.

DR. NEVO: Thank you. Maybe we'll get a chance for the authors, if you want to say one last word or respond to any of the comments? Kate, you expressed an interest. Maybe all three of you can come up here because we might get questions.

DR. GENTZKOW: I don't have anything to say other than those are great suggestions. Thank you.

DR. HO: I had a couple of things I wanted to say. One is in response to the comment on the last slide. You said you would expect discounts to be lower for high capitation insurers because hospitals are bearing risk.

In fact, hospitals aren't bearing risk here. It's the physicians that are bearing risk. Hospitals are bearing no risk, so I didn't think that that was such a counterintuitive finding, so I think that we're okay on that front.

You said we need standard errors, we definitely need standard errors. We'll get there. We haven't gotten there yet, and there was a comment at the beginning about whether consumer preferences should be allowed to differ across consumers, across patients, and we do some of that.

So we allow preferences for quality to differ across severity groups which is important. We also allow average preferences for quality to differ across insurers, so that we're allowing for selection of different types of consumers into different insurers, for example.

What we're not allowing for is for consumer preferences for price to differ across types of insurers. We don't think that's likely as we explained in the paper. We also have a test for that in the inequalities or some kind of a test at least where we estimate the inequalities separately for sick versus less sick patients and find essentially the same price coefficients, but other than that, they were very helpful comments. Thank you.

DR. MILLER: I think Allan was right on. Thank you.

DR. NEVO: Questions?

UNIDENTIFIED SPEAKER: Thanks. This is more an invitation to clarify I hope. Nate, I think near the beginning you said something like you assume that where the data differ from the predictions of the model, it's because of incorrect data measurement.

DR. MILLER: That's right.

UNIDENTIFIED SPEAKER: If taken out of context or perhaps even if taken in context, that might sound like a rather aggressive assumption. Do you want to clarify what you're really assuming there?

DR. MILLER: Do you mean that it might be data due to incorrect specification? So in general once you aggregate out the predictions of the model and compare that to the data, they're not going to be exactly right, and one needs to talk about what that error is and how to deal with it.

I guess if you're willing to say the error is endogenous to the plant locations and to the cost and demand shifters, and the error here being the discrepancy between the predictions and the data, then you get the nonlinear multiple equations, nonlinear least squares.

If there's some relationship between that error and the plant locations, which I kind of scratch my head on why that might be the case, but if it happened, what one would need to do is instrument, and you would get a slightly more complicated estimator, but I think estimation would still be feasible.

Essentially what one would need is an instrument that is correlated with the equilibrium predictions, that's not correlated with whatever the term would be. You would have to think hard about that what actually is in the near term.

DR. NEVO: Other questions?

UNIDENTIFIED SPEAKER: I have a question for Kate. We heard this morning from Roman Inderst about some of the agency problems that exist with financial advising. It would seem that the introduction of any incentives to get doctors to make decisions which are lowering costs and maybe in the interest of the HMO or whatever is likely to be -- well, I guess I worry about that that introduces an agency problem with the patients.

You started out the paper, motivating it by saying we're looking for ways to reduce costs without reducing quality, and somehow that seems to me to be impossible to be able to do, and then I worry about this potentially agency problem and should there be some transparency? Should patients be aware that their doctors have these capitation things going on? And is there anything in the work that you've done that would speak to that?

DR. HO: That's a great question, and it's something we haven't looked at in detail so far. We've looked at women going into the hospital to give birth, and in some sense we're less worried about quality there, the quality of the hospital they're going to than we might be if these were patients with cancer or some more severe illness. We have plans in the works to try to understand the trade off being made here between quality and price.

So far we're just looking at whether price matters and the extent to which it matters. There's obviously a follow-up question about how much are we losing in terms of quality.

From the initial analysis we've done, it actually looks as if patients are going to cheaper hospitals. They're cheaper because for reasons that might not be related to quality. They're cheaper because they're hospitals that aren't very high tech. They don't offer transplants, for example. Women giving birth don't care if the hospital offers transplants or not.

It's not about the number of nurses in the hospital. It's not about C section rates, so our initial analysis is saying quality isn't suffering to a large degree, but I agree that there's a lot we could look at there that we haven't done.

UNIDENTIFIED SPEAKER: Thank you. My question is for Matt in the first paper. I tend to think about the extensive preferences, but even I was surprised by this drop of 60 percent once you moved, so I was wondering what would be the explanation for that? And maybe it's just a question, I don't know how you're measuring migration of the household because if only one part of the household, one member of the household migrated, you might explain the -- if 60 percent of people came from California to marry a women in Washington, you might expect a jump in broad terms.

So I wonder if you can differentiate if the whole household moved or only one part of the household.

DR. GENTZKOW: Good question. In response to the second part, we actually have data at the individual level within household. We had individuals fill out the survey. In the large share of households if there are multiple people, they both filled it out. We did something pretty simple with that where we just select whoever does most of the shopping and used them.

In the panel, household composition is remaining constant over that time so I think it's not an issue. In terms of what explains the jump, which I think is really an interesting question, we have some descriptive evidence in the paper, prices. This you really could have really expected might go either way, but prices are lower where brand shares are high.

So when you move from a state where Folgers is popular to a state where Maxwell House is popular, you're getting lower price for Maxwell House. That's one obvious thing that would explain an immediate change.

Second, something that's harder to measure but that I think intuitively we think all of the literature on how stores allocate shelf space says to a first approximation, they should do that proportional to market shares, so you move to somewhere where Maxwell House is popular and if you are the kind of person who just walks into the store and picks whatever is there, you're going to immediately start buying more

Maxwell House.

Advertising and promotional activity are also positively correlated with this. We can kind of do a decomposition exercise in terms of the covariances. How important those different things are, we can't really separate them, but I think the behavior is consistent with that, and we're doing some follow-up work thinking about the supply side and what explains how firms' decisions will be different in a place where those things are higher.

Those supply side variables are what make this persistent over long periods of time because it means if the old people in this place like Maxwell House, if that makes it optimal for me to charge lower prices and have more variability, then new consumers coming into the market will learn to like Maxwell House as well and it's going to persist over generations.

DR. NEVO: Are we done in terms of time? Well, let's all thank the speakers and discussants.

(Whereupon, a brief recess was taken.)

DR. ADAMS: We have lunch out on the tables, and I should say that our calorie information is exactly the same as our price information. Please come back here at about ten past 1:00. Feel free to bring your sandwiches back in but not your conversation.

(Whereupon, at 12:40 p.m., a lunch recess was taken.) AFTERNOON SESSION (1:10 p.m.)

DR. BECKER: I'm going to make a quick announcement as people are filing in. We've made a change in how we are doing wireless in this room, so if you would like wireless access here, you need to go out to the table in the front where you first signed in. There's a list. You can put your name on the list. There's also someone out there so you can ask them if there's any confusion about it.

It's an honor for me to introduce an economist who has been incredibly influential in both how we do things at the FTC or how we do things in Consumer Protection, but also how economists everywhere think about economic decision making.

David Laibson is the professor of economics at Harvard University. He's advanced economics by looking inside the black box in a lot of areas including intertemporal tradeoffs and decision making under limited information.

His research has been published extensively both in economics journals such as the AER and QJE, but also in journals outside of economics such as Science and the Journal of Neuroscience.

His contributions have been so extensive that

I feel he deserves a long introduction, but I'm going to turn the floor over to David Laibson.

(Applause.)

DR. LAIBSON: Thanks a lot for this opportunity and the invitation. It's great to be here. I want to tell you about some work with Sumit Agarwal, who is at the Fed in Chicago, John Driscoll, who is at the Fed here, and Xavier Gabaix who is at NYU. And as you probably all know, this is not in any way reflective of the views of the Fed.

So this is a talk about financial decision making over the life cycle, and it's just a motivating example. Considered Brooke Astor. This is shortly after her wedding to the Astor family at age 51, and you know the story of her life cycle.

She, for the next 50 years after that, becomes the leading member of society in New York City, the major advocate for almost every charity in New York City. She eventually receives the highest civilian honor available in this country, the Presidential Metal of Freedom.

Shortly thereafter she begins to decline cognitively, eventually gets an Alzheimer's diagnosis, and is then the victim of psychological and physical abuse from her son. Eventually he's convicted of grand larceny and is now in jail. She died around 2006, I think. I'll be talking about maybe less extreme examples of those kinds of sad endings.

So there are a lots of performance peaks wherever you look in economic life and in life more generally. Here are a few examples. You want to find a good dictator. He or she should be about 45. You want to find a good economist. There's debate actually on what the right age is for great economic research is. I want to go into that.

Today we're going to talk about financial performance, how people make financial decisions in the domain of credit card markets or in the area of credit markets, and I'll talk about performance in ten different areas. We'll basically find that performance rises and declines with age in the cross-section and we'll be measuring performance based on fees and negotiated interest rates in loans.

So, these are the markets that we studied. I'm going to go through them one at a time later, but we basically got data from a bank that shared with us every bit of information they had in all of these markets, and we'll show you what we learned. We had all the data that they had on their borrowers.

Now, when we talk about this pattern of

rising performance and then falling performance, there are obviously many stories that might come to mind. That pattern was observed in this data, primarily in the cross-section. So, when we talk about reduction effects or current effects perhaps driving that pattern, and I'll talk quickly about why we think that's not what's going on here.

There's a large and, I guess, small but rapidly growing literature beginning to think about how cognitive performance affects economic decision making in the domain of age. I think some of the key contributors are Korniotis and Kumar and Zinman, who is here today. There's a lot of literature thinking about how differences in cognitive capabilities affect important economic outcomes.

So I will present ten different credit markets and talk about behaviors in those markets and then to discuss quickly the various explanations, and I'll emphasize age related effects as opposed to overt effects or selection effects.

This is the first set of markets we want to talk about. It's loans collateralized by a home, so we'll be talking about home equity loans and home equity credit lines. This is again proprietary data from a single bank, and everything that I show you comes from the same bank who wishes to remain anonymous.

The data for homes involves 75,000 contracts, and these are contracts from 2002. Again we observed everything the bank observes, and we're going to put all those data, all those characteristics on the right-hand side of these equations. The key variable that we will studying on is the age line which will go on the right-hand side of the equation.

So let's take a look at how interest rates, home equity lines vary with the borrower's age, controlling for every bit of information that the bank has about the borrower and keeping in mind that banks can't make age contingent interest rates. That would be illegal. Here's what we see in the data.

You can see this U-shaped pattern. If you're a young borrower, you're going to pay, in this, example about 6.4 percent. If you're a middle age borrower, you're going to pay 5.4 percent, and if you're an older borrower, here going up to age 80, you're going to pay about 6 percent.

There are no standard errors in any slide I'm going to show you because they're all tiny, so tiny that you just see little parallel lines walking this thing down and up because we have so much data.

Here's the same plot now for home equity credit lines, and we see basically the same pattern here as well. We see a hundred basis point improvement as we go from young to middle age, and then a 75 basis point worsening as we go from middle aged to age 80. I'm going to keep moving. We're going to see ten markets in total.

In the next market, this is what we call reactive behavior. Back in the day before the Fed banned it not so long ago, credit card companies engaged in the following routine: They would say to a new client, "Please transfer balances from your old credit card company to this new credit card company, us," and the big letters would say, "And you're going to get a low interest rate on the transferred balances."

Then the fine print would say that every time you make a payment on the new card, we'll be crediting your balance transfer first before crediting the actual new charge, which means that as you make new charges, you're accumulating high interest rate debt and effectively with each payment paying off the low interest rate debt.

By implication the optimal strategy is to do the balance transfer and then put the card in a desk drawer and forget about it until the low interest rate period expires. Not using the card is the only way to take maximal advantage of the transferred balances at a low interest rate.

Now, that's pretty hard to understand. I didn't get it the first ten times I read these inducements back eight or nine years ago when they were popular. Let's see who gets it in terms of the people who are actually taking up these offers.

So take a look first at the light blue greenish line. So it's like a blueish, greenish, grayish line. And what that's plotting is the fraction of individuals who get one of these balance transfer offers and don't get it because they keep using the balance transfer card to make new charges, basically effectively losing the ability to fully take advantage of the balance transfer.

You can see that a lot of the older individuals who are bucketing over 65, over half of them, don't ever get it, meaning they keep using the card for the entire duration of the low interest period. Among middle aged borrowers, only 25 percent never get it, and among young borrowers, 45 percent never get it.

And the opposite of that are those that get it right from day one and never use the card -- I'm actually amazed at how many people fall into this category; they're far brighter than I am -- 25 percent get it from day one and don't use the card in the youngest age bucket. 45 percent get it from day one and don't use the card in the middle age bucket and about 20 percent get it from day one and don't use the card from the oldest aged bucket.

Now, I'm going to go through seven more categories of credit card or credit market behavior. We'll talk about fees and then we'll talk about a bunch of interest rates.

So first fees: You can see here the frequency of late payment fees, the dashed blue line up here, the frequency of cash advance fees. This is a monthly basis and the frequency of over limit fees down here, and again you see the same U-shaped pattern, though the magnitudes now are less pronounced since the scale is somewhat compressed.

Here we see auto loan interest rates and again we're on the right-hand side controlling for everything that the bank sees. You can see here the same U-shaped pattern, though a bit less pronounced. Now the differences here are from 9 percent down to about 8.7 percent, 30 basis points, and then maybe 20 basis points back up on the other side.

Here are credit card APRs by borrower age. We see a robust decline in the beginning of life and then a nature or slightly rising pattern later in life. Here are mortgage APRs by borrower age. Here again you see a decline and rising later in life. This is about 40 basis points different here. Here you see small business credit card APRs, again the same U-shaped pattern.

Now, if you take all that data, and you push it all together and you ask, "Where are individuals getting the best deals, paying the lowest risk adjusted interest rates or FICA adjusted interest rates?" It turns out it's about age 53. We estimate that by taking the middle aged population and fitting a quadratic to that population, so it looks like the peak of financial performance, those who are doing best in the data, are 53 years old or thereabouts.

There are a lot of explanations that are plausible candidates for this pattern. I want to focus on the age related effects and then try to talk you out of selection effects and cohort effects.

Let's begin with age related effects. We'll begin with little background. There are basically two kinds of intelligence that we humans seem to have. At least psychologists bucket it this way. There's crystallized intelligence, and there's fluid intelligence. Crystallized intelligence is the name of the vice president. Is that a hard one? Believe it or not, that is a hard one for many Americans.

Fluid intelligence is your ability to see a new problem and solve it, so crystallized means a familiar piece of information. You've heard it before. Can you basically recall it? Fluid intelligence is your ability to confront a brand new problem and get it right, so let's take a look at fluid intelligence tests.

This is the bad news, so I'm about to give you some awful news, which I hope you'll accept rather than deny. These are tests for fluid intelligence. Here's a memory test. Here's a list of ten words, which of them can you remember and then write down after the list disappears? Here's a spatial visualization test. Look at this two-dimensional object on the left, cut it in your mind away from the background paper, fold it accordingly, and which object on the right will you have reproduced?

Here's a matrix reasoning. That's right, it's not so easy. Here's a matrix reasoning task: Which is the missing object in the lower right-hand box. And then finally perceptual speed, take a look at those two objects in each row and very quickly tell me whether they're the same or different.

So these are examples of the kinds of tests that are used to measure fluid intelligence.

Now, unfortunately in the cross-section, this performance declines very, very sharply over the life course, so if you look at 20 year olds on fluid intelligence, this is on any measure, perceptual speed, 20 year olds perform at about the 73rd percentile of the adult population while the average 80 year old performs at about the 16th percentile of the adult population.

Here's data from the HRS, which is not in a cross-section. This is data that we're actually controlling for fixed effects because people are given the same questions repeatedly. And you can see here, for example the answer to the question, if the chance of getting a disease is 10 percent, how many people out of 1,000 would be expected to get the disease.

A fraction of people who answer a hundred; 80 percent at age 50, not so great; 50 percent at age 90, not so great, too. Here's another question. Can you divide two million by five to get the answer 400,000? At age 50, about 50 percent can do that. At age 90, about 10 percent can do that.

Now, part of the story is dementia, more bad news. Basically the prevalence of dementia doubles every five years with age, so it goes from .8 to 1.7 to 3.3 to 6.5 to 12.8 to 30.1. Every five years of your age, your likelihood of having dementia doubles.

Moreover, there's the risk of not quite having dementia, but being on the door stop of dementia because you're cognitively impaired but you're a few years away from a full-blown dementia diagnosis. It turns out in the 70s, that's about 16 percent of the population. In the 80s, it's about 29 percent of the population, and in the 90s, it's about 39 percent of the population.

Put this together and you will see that approximately half of adults between age 80 and 90 either have cognitive impairment, just short of dementia or full-blown dementia, half the adult population in that age band.

So what is average dementia rating between age 80 and 90? It's about a diagnosis halfway between mild dementia and moderate dementia. So put all this together; what do we see? We see a pattern whereby crystallized intelligence, people's ability to remember that the vice president is Biden, is basically rising dramatically over the life course as they gain more experience, but fluid intelligence, the ability to solve a new problem is rapidly declining over the life course.

If you think about our decisions as being in some sense a composite of experience and fluid intelligence, it's not really surprising that's going to rise and then fall over the life course, so performance would be peaking around mid age.

One test of the hypothesis is to ask how performance changes as people enter different activities at different points in life. And the prediction would be that when you enter an activity, there's a period of rising performance, as you gain experience. That's that concave, crystallized intelligence function, being hit by declining fluid intelligence as you get older and older.

So you would expect performance peaks to occur later, and later the later in life one begins a new activity. That's exactly what we see when we look across the ten domains of behavior that we've been studying. Performance peaks are later as people begin some activity later in life.

Now, another possible explanation for all of this is a cohort effect or a selection effect. Let me first turn to selection effect. So there are two kinds of selection that we're able to rule out. One kind of selection is that individuals who are in our sample are unrepresentative of the full population at those different ages, so perhaps the young people in our sample, because we're dealing with again selected individuals in the database of this large bank. And maybe the old people in our sample are not representative of typical young and old people in terms of education and other characteristics.

We can go to the SCF and ask whether debt holders are more or less sophisticated across the age range and in fact this works just the opposite way. Young debt holders and old debt holders are in fact more sophisticated than the average individual based on education and income in those age buckets. So, the first selection effect actually goes against us. It's not pushing results in our direction.

The second effect is that perhaps middle aged borrowers are just different, less risky than the young and old that borrow. There again, the evidence points in exactly the opposite direction. If you look at the data that we have, the middle aged borrowers are in fact the most likely to default and the old borrowers and the young borrowers are in fact the least likely to default, so none of those results are explained by the default rates or the risk characteristics of those in middle age.

We don't think cohort effects are driving the data for many, many, many reasons, but the most important is the last one on the slide which is when we go back in time to data ten years earlier, we see the exact same pattern of performance peaking at age 53.

If it were a cohort effect, we would expect to see that peak shifting as they allow the sampling for the cross-section, but we don't see that shift. We see the peak of performance at 53 staying in the same location regardless of the time period in which the cross-section is sampled, so that seems to rule out cohort effects.

Now, there's cost of the time effects, but they go in exactly the opposite direction. If you think that it's the cost of time that's varying over the life course, you would expect that older adults who are retired and have lots of time would be making the best decisions and young adults who were presumably not facing a lot of time demands because they have low wages and small families would also be making good decisions, but that's not what we see.

It's the middle aged that make the best

decisions, despite the fact that they have the most demands on their time and the highest opportunity cost of time, so we don't think that the cost of time can even plausibly explain any of these effects. As I mentioned before, it's not default behavior since the default rates in our sample of prime borrowers are highest for the middle aged and lowest for the young and the old.

So to conclude, we see a robust pattern of low performance for the young and the old relative to high performance for the middle aged across all the markets for which we have data. There are others who have also looked at their data and found similar patterns.

I believe Fiona did that at some point. And there are several other collaborative teams who have been looking at their own data control. And they're seeing the similar pattern that when you control for the characteristics that are relevant, say FICA scores, say loan to value ratios, you find that the middle aged are performing better on economic tasks than the young and the old.

Now, I want to emphasize one policy issue before finally concluding. The paper that I'm discussing here actually contains a large range of policy observations, but I don't have enough time for that today. We actually have a remarkably perverse regulatory framework, vis-a-vis financial decision making in this country.

We have an excellent regulatory safety net for middle aged people in the form of pension plans that are regulated by ERISA, and then because of that regulation an investment committee that acts as a fiduciary that screens out bad choices, that basically creates an incredibly safe sandbox for employees between age 20 and retirement, 65.

Then after you retire, you exit your retirement savings plan. You exit your pension plan, and you end up in an IRA rollover that has no fiduciary protection, that has absolutely nothing that is even approximate to ERISA. It's basically a Wild West, so we've completely reversed the appropriate regulatory environment.

Instead of providing regulation that really protects vulnerable older adults with large chunks of retirement wealth, we have built a system that protects the middle aged, according to our results, those who need it the least.

Let me summarize. Older adults experience substantial declines in analytic cognitive function, declines that, I think, we should be very worried about as regulators. And certainly we should be thinking about as economists because that can lead to all sorts of interesting economic behavior that doesn't jive with the rational actor model. The data we've looked at today, data from credit markets indicates that the middle aged appear to be doing relatively well on every dimension that we were able to study relative to the young and the old.

There are a lot of open questions. This is just, I think, I hope, the beginning of a research program to come covering questions like how important are these losses, questions like to what extent do individuals anticipate these changes in their cognition as they approach retirement and prepare for them.

There's lot of things you can do: irrevocable trust, bringing powers of attorney. Are individuals able to avoid these problems by creating institutional protections for themselves either with the help of the government or just on their own with the help of some kind of legal intervention?

Does financial education help? Do third parties help? Can we delegate successfully these decisions or are we stumped at making them ourselves and thereby making them badly? How is the market responding to this? Is the market creating a solution or is the market evolving to exploit these vulnerable older adults and creating opportunities to separate them from their money?

Then finally: What is the appropriate regulatory response? I believe at the very least we should level the playing field and give older adults the same protections that we've been offering middle aged adults to date.

Thank you.

(Applause.)

DR. LAIBSON: So I don't know whether we want to take questions now? I am happy to, but I understand if you want to get back on schedule.

DR. ROTHSTEIN: We can take questions for five minutes.

DR. SCOTT MORTON: Fiona Scott Morton. I just want to ask a question. You said you put everything that the bank knows on the right-hand side. Why then is there a residual? What is that residual?

DR. LAIBSON: With regard to interest rates, I go into a bank and I negotiate for an interest rate.

Now, it's true if I'm just reading off a schedule of interest rates, there's no residual, but that's not the way it works. There are all sorts of steps in which I can end up going off the standard schedule and ending up asking for a better rate, walking out if I can't get that rate, et cetera.

We think that's the key element, and we actually have evidence of where the negotiation is breaking down. It turns out that the middle aged have a much better understanding of the value of their homes than do the young and old. So when you come into the negotiation with a very poor understanding of the value of your home that gives the bank an opportunity to deviate from the standard routine and offer an interest rate that is basically punitive.

That's the case at this bank; we spoke to the origination group, and that's the way it works.

UNIDENTIFIED SPEAKER: David, I'm sure you thought of this: Is income an alternative explanation here? The greater your income is, the more flexibility you have to be able to take advantage of the optimal things that come along or the less constrained you are to be able to make better choices and income peaks around 53 as well?

DR. LAIBSON: For most of these data sets, we have income, so that is an argument on the right-hand side. So, I don't think that would be it, though I guess it's one issue that I thought of as you were asking the question. We have probably linear and log income on the right-hand side and maybe there's some interesting alternative function that might be better at soaking up some of that variation but we do have log income on the right-hand side.

DR. GENTZKOW: You sort of dodge the question of magnitudes here by saying that's an open question. It seems like there's a lot of direct stuff in your data about how big these things are and many of them look small kind of eyeballing interest rates, but what do you know about the size?

DR. LAIBSON: So we think the effects are going to be absolutely enormous when you add them up across every single domain. If you're going domain by domain and you ask how much more interest are people paying at age 85 on their credit cards? A couple hundred bucks a year, but you then think about the credit card and the mortgage and the low returns on their financial products.

So Korniotis and Kumar estimate a 200 basis point lower risk adjusted return for 80 years old relative to people in the middle age, and that's a terrifying number. If it's 200 basis points and that's your retirement wealth and you're talking about let's say a $500,000 pool of assets for an upper middle class family, that would be, I believe, $10,000 a year in differential returns, presumably just as a function of your age.

That's not even taking into account things like fraud, so there's this enormous concern now in Washington that there's a huge set of market actors targeting these individuals and stripping them of their wealth in ways that are borderline legal. And that's another huge category of losses, so no one's added it all up, but if we think if one did add it all up, it could be that you're stripping from a typical 80 year old $5,000 to $20,000 a year by the time you've taken all the categories and strung them together.

DR. GENTZKOW: It would be useful to know those.

DR. LAIBSON: Yes, and we do have those numbers in the paper. Within the domain, you're right. It's on the order of $500 a year within this domain.

UNIDENTIFIED SPEAKER: So given that the bank's data may not perfectly capture liquidity constraints or the consumers expected near term or medium term liquidity constraints, have you thought about whether you might be just picking up something about the life cycle of credit demand on the loan pricing stuff?

Is the early 50s the age at which liquidity constraints bottom out because kids are leaving the house and you're not getting hit with uninsured medical expenses yet and so on and so forth, and so maybe this is sort of where demand for credit bottoms out, so people bargain hard, and if they don't get the deal they want, they walk away?

DR. LAIBSON: Well, I'm skeptical about that story for two reasons. The first is that if it were the kind of life cycle of liquidity needs, we usually think about liquidity needs being the lowest just before retirement, when you have this big pool of assets that you've saved up for retirement, whereas we think about household and the 50s as still in the process of paying college tuition and stuff like that.

So I would have thought that the liquidity story would have generated the best rates for 64 year olds who have the big pool of assets. But the other thing that makes me think it's not liquidity is that we have got a lot of characteristics on the right-hand side that are basically measures of the financial desperation of the household, like FICA scores and like loan to value ratios.

So again in some sense those controls are in there, though they're imperfect.

UNIDENTIFIED SPEAKER: I don't know whether this is in the data you have, but if we think that part of it might be different bargaining behavior, it would be interesting to know what fraction, by age of initial negotiations, do not lead to a final transaction.

DR. LAIBSON: We don't have that data and I should emphasize that would be great data to get. It's not in our data set. We wanted that and couldn't get it out of the bank. When I say negotiation, I actually mean two things that fall under that category.

The first is getting a bank to give you a better offer, and the second is walking out the door when the bank gives you a bad offer. So, you could think about some people just coming in and getting something that's inflexible and taking it, whereas the smart folks say, "That's not a competitive offer, I'm out of here." I think both of those margins are active in terms of this negotiate premium.

UNIDENTIFIED SPEAKER: It's fascinating data and I think it really speaks to our concerns in consumer protection about basically which groups of consumers we have to protect, the very young ones and the very old ones.

We see it in many different markets as well, but now you have started talking a lot about social interactions which may explain some of the outcomes, the negotiations, and actually you say here you have data, but then you tell a story.

The story you basically tell is one of cognitive ability and knowledge. Maybe some of your findings can also tell a different story, and it's a story that's told in other markets. For instance, in your case, Office of Fair Trading has done a study of doorstep selling.

They found obviously that for those consumer groups which are particularly vulnerable, particularly the very old ones, their story was not one of cognitive abilities, but one of social inference, for instance, like which kind of people are not giving into pressure, which kind of people are basically falling in my topic advice and recommendations. And from this perspective, I think you can also tell a story that would explain some of your data.

From some regressions I've been running, we have seen that those were not naive with respect to advice. It seems to be those which are not in a business situation at the moment, for the young ones and the very old ones, whereas those say you are 40, 50 years old, these people may think about each and every transaction from a business oriented sense, so maybe that can also explain it.

That's like different concepts like reasoning. These may basically think about these decisions when they make them together with the bank official, like a social interaction flavored with a cup of coffee. Whereas the 40 or 50 years old will think of business savvy, so maybe that can also explain it. But the nice thing you have in here are different decisions, those which are taken alone and those that are taken in the social context. Maybe you could try to do more on this.

DR. LAIBSON: So that would explain some of our settings, but as you just anticipated I think in the last sentence, some of our settings are completely private settings, meaning there's no one on the other side of the decision, for example, the eureka moments. And I think the fact that we see the robust pattern across all the settings, those that have a social possible interpretation and those that don't, makes me think that while the social story is plausible, it can only be part of the story.

Thank you very much.

(Applause.)

PAPER SESSION TWO: INATTENTIVE CONSUMERS

DAVID LAIBSON, Chairman, Harvard University

PRESENTER: MICHAEL GRUBB, Massachusetts Institute of Technology, Sloan School of Management

DISCUSSANT: Ginger Jin, University of Maryland

PRESENTER: JONATHAN ZINMAN, Dartmouth College

DISCUSSANT: Karen Pence, Federal Reserve Board

PRESENTER: NICOLA LACETERA, University of Toronto

DISCUSSANT: Kory Kroft, Yale School of Management

DR. LAIBSON: I'm now going to run this session, but I'm not going to sit there because I want to be able to see what's going on. So we've got three terrific papers and the organizing principle for this session which kind of emerged by accident is inattentive consumers. All the papers speak to that issue, and our first presenter is Michael Grubb from MIT.

DR. GRUBB: So thank you very much for having me. I'm going to talk about penalty pricing and regulation requiring firms to disclose at the point of sale whether penalty fees might apply on that transaction.

The motivation for this research comes from three observations. The first is that there's a lot of situations where services are priced non linearly for which a consumer may be fully aware of the contract terms and yet unaware of what the marginal price of any particular transaction is.

So, for example, imagine a cellular phone customer. They're fully aware that their contract marginal price of minutes is zero up to their 500 minute allowance. Thereafter, it's 35 cents a minute, but if they haven't kept track of how much they've talked on the phone or it's a family plan and they haven't heard from their spouse how much their spouse used, they don't know whether or not they're over that 500 minute limit. And then the phone rings, they don't know whether the marginal price is zero or 35 cents a minute.

Similarly, a bank customer might be fully aware that a debit card swipe transaction has a zero transaction fee. If there's no money in their account, the fee is a $35 overdraft charge. And yet, if they haven't kept track of how much money they're spending, they might not know whether there's money in their account, and therefore not be sure whether the next transaction fee is zero or $35.

There's growing empirical evidence showing that people are in fact uncertain about whether these penalty fees apply at the point of sale. We're going to hear some very interesting stuff on that from Jon

Zinman next.

The second point is that this unawareness of the marginal price and point of sale is endogenous so firms have the option to could make your phone screen flash bright red and say, "Overage fee applies" when the phone rings to let you know before you answer. They don't. You can go online and find out how many minutes you've used or call a number, but it's not put right in front of your face.

Similarly, you can check your checking account balance online or do a balance inquiry at the ATM, but when you swipe at Starbucks for your coffee, it won't say that an overdraft fee is about to apply and ask if you would like to continue.

So my question is: What would be the effect of requiring firms to make this disclosure of whether or not a penalty fee applies to that transaction and would it be a good idea?

Related to the second point - and this may not be news to people here -- there's been some recent regulatory attention on this issue. So for example, for these particular two applications of cell phone charges and overdraft fees (We'll hear a lot more about the overdraft fees in a minute.) since July, the Fed is requiring banks to actually ask customers in advance if they would like the overdraft protection service.

With respect to cell phone charges, there's recently been a bill shock regulation in the EU, and the FCC is considering something similar here, which would require firms to send a text message to consumers if they start triggering these high fees, and their bill starts increasing, so that's the motivation.

What I'm going to try to do, if I can manage it in the 20 minutes, is first talk about how consumers make choices, at least as I model it, when they're unaware of the price. Essentially the answer is going to be they're going to respond to an expected marginal price rather than what it actually is, and then I'm going to go through three models.

The first model is a benchmark model where everything is as simple as possible. There the main result is an equivalence result. Essentially it says that I'm thinking of two different types of consumers, ones who keep track of your usage -- I am calling those attentive consumers -- and ones who don't keep track of their usage, so unaware of whether or not a penalty fee applies. I'm calling those inattentive consumers.

The first result is it doesn't matter whether consumers are attentive or inattentive, and if they're inattentive, it doesn't matter whether or not we require firms to disclose information about the price then. Profits, market shares, consumer surplus, allocations are all going to be unaffected. The only thing that might differ is actually the prices we observe.

The second model I'm going to talk about enriches things in a way where I think is particularly applicable to the seller of phone context.

So, companies don't just offer one contract; they offer a menu of contracts. Why is that? Because consumers are different. Some expect to talk a lot on the phone. They might buy a plan with a high monthly fee, but a lot of included minutes. Others expect to talk very little on the phone and pick a cheaper plan.

I want to introduce this sort of heterogeneity into the model, and there's a model for price discrimination. Now I find that the issues of inattention and disclosure actually matter a lot.

One of the things I find is that now it definitely makes sense for firms to charge penalty fees and to make them surprise fees. And so endogenously for them to decide not to disclose at the point of sale whether a penalty fee applies, but interestingly, regulation requiring them to do so can be counterproductive, and will be counterproductive, if you believe all the assumptions in the model, in fairly competitive markets.

Now, I think this might apply to the cell phone pricing and so you should be cautious about this new FCC bill shock regulation. I think it doesn't apply as well to the bank overdraft charges case because although banks offer different types of checking accounts with different terms and fees and tries to sort different target customer groups into those different accounts, as far as I'm aware, in the past they haven't used the overdraft fees to help in that sorting. The overdraft fee structures have been similar across all the types of accounts.

So, this price discrimination doesn't really apply. I have a third model that I think may be insightful to the overdraft fee case. My assumption here is that consumers might underestimate their demand for the service, in this case essentially underestimate how much they're spending, and so underestimate the likelihood of going below a zero balance and triggering overdraft fees.

There again I find that it makes sense for firms to charge penalty fees and not tell consumers at the point of sale that they're about to apply. Here I don't have clear conclusions about whether price posting regulation would be good or bad for welfare. It could go either way, but I think there's potentially a strong benefit of the regulation protecting consumers from exploitation, and this is true even in a competitive context.

So here's the starting model, my starting point. It starts out at time zero. That's a contracting phase. At this point, differentiated firms have a chance to offer a nonlinear contract, which I will explain in a second. Consumers decide to sign up or pick an outside option. Two subsequent time periods, consumers make a perfect decision so this is to try to make this as simple as possible.

These purchase decisions are binary -- buy or not buy -- so I pick a quantity that's either zero or one after I learn my value for the object. -So, if you think of this as cell phones, I have the possibility of making zero, one or two phone calls on the billing cycle. At time one, I learned my value for the phone call is V 1 so I decide: Do I make it or not?

Now I can make sense of what this contract is. The contract essentially would offer some fixed fee and plus a marginal charge for the total usage, plus a penalty fee if you buy both units, so I think of this penalty fee as applying. If you buy both units, then potentially there's an additional marginal charge which I'm calling a penalty fee.

I have some standard risk neutral payoffs here. Consumer utility has an additive brand shock that's going to allow me, for instance, to think about a Hotelling duopoly where firms have some market power and firms have constant marginal costs.

So if consumers were attentive, that would be the end of the story. I'm going to allow consumers to be inattentive, which I'm modeling as having imperfect recall, so at time two, when they're making their second buy or not buy decision, they can't remember whether or not they purchased in the first period. And hence they can't condition their choices to buy in the second period on whether or not they bought in the first period.

In this case, the optimal strategy is to buy, even if your value is above some threshold V star, where the optimal threshold is the expected marginal price. So, I know I'm always going to have to pay at least the base marginal fee P, but if I buy in the other period, that means I trigger the penalty fee and then I also have to pay the penalty fee. At time two I can't remember if I bought yesterday, so I have to think, "Well, what's the probability?" Well, the probability is that my value was above the threshold last period.

Again in the first period there's no memory problem, but I don't know what's going to happen in the future. What's the probability I'll buy in the second period again? My value is above threshold.

The primary policy intervention I'm going to consider is what I'm calling price posting regulation. It's a disclosure requirement. It's the requirement that the firm tell you whether this penalty fee applies. Essentially because there's only two periods, this is equivalent to disclosing your entire past purchase history so totally solving your memory problem that you can't remember how much you've used the service or how much money is in your account.

Another intervention I'm not going to focus on in the talk is banning penalty fees, essentially requiring firms to charge a constant marginal price and offer a menu of two part tariffs. That's going to have the same quality effects as the price posting regulation.

So the first main result in this benchmark model is the equivalence result I mentioned in the beginning, and this says if consumers are all the same, they draw these values, VTs, from some distribution, that that's the same for all customers, and they correctly understand their value for their product. They correctly understand this distribution F. Then inattention and price posting regulation have no substantive effect.

So welfare, profits, consumer surplus, and market shares, are all unaffected, and allocations are first best, whether or not there is disclosure, whether or not consumers are inattentive.

In the attentive case, the only prices we would see in equilibrium are marginal cost pricing, so firms will set marginal price equal to marginal cost to get the efficient surplus and extract as much as they can of that as profits through fixed fees. If people are inattentive, we might see different prices.

The only prediction of the model is that prices will be set so that the expected marginal price, the structure that people make to purchase decisions, is equal to marginal costs. So, we're still getting first best allocations in surplus, but that could now be implemented with a variety of different prices, including a zero base marginal charge, and the only marginal charge being a penalty fee that's sufficiently high that an expectation of marginal price is equal to marginal cost.

I think this is the result that Jamie Dimon should have had in his back pocket when he made the following argument. He said: "If you're a restaurant and you can't charge for the soda, you're going to charge more for the burger. Over time, it will all be repriced into the business."

So that's what this proposition is saying is that if we have something like the bill shock regulation, what it might mean is that we get rid of steep penalty fees, but it's going to be made up by raising marginal prices elsewhere, and the things that we care about like profits and surplus are not going to be affected.

Now, I don't actually think it's the case that this doesn't matter. I think the equivalence result depends importantly on all of the things that are missing in that benchmark model. The first thing I want to add is this heterogeneity, an incentive for firms to price discriminate between different groups. So, how does the model change?

Well, now at the contracting phase, firms are not going to offer just one contract, they're going to offer two contracts, a low contract and a high contract. Why is this? Because there are two types of consumers, those that are receiving a low signal and a high signal. What's the difference between these two groups? Well, at times when they're deciding, "Should I make the call or not" or "Should I make the transaction or not," the value draws that they're getting are coming from different distributions, from either a high distribution or a low distribution. So high consumers are ones who have higher values to the product and so they're more likely to buy.

Here I'm going to focus on the case where I have two duopolists who are located at opposite ends of the Hotelling line and consumers uniformly distributed along it. There again I have different transportation costs for high consumers and low consumers. Think of the high consumers, they're willing to pay more to purchase to make phone calls. They're a high income group. They're also more willing to pay to go to their preferred brands. They have higher transportation costs, H, and the lower consumers have lower transportation costs, L.

The first result familiar from standard price discrimination models would be if the consumers are attentive, they can keep track of their usage, and they know what the marginal price is at any point, firms are going to offer contracts that are going to include penalty fees. In all, equilibria allocations are going to be inefficient.

The high type is going to get a contract with marginal cost pricing and make the efficient allocation first best. The low type is going to choose a lower fixed fee contract with marginal prices above marginal costs, and their allocation is going to be distorted downwards, so there's necessarily going to be an inefficiency there.

What's interesting is when consumers are inattentive, we needn't have that inefficiency, and we won't in a fairly competitive market. So I want to keep these transportation costs positive if they are actually equal to zero. If we had perfect competition, then pricing would just be at cost. There would be no penalty fees. There would be no scope for disclosure, so I allow it to be positive so there's some market power, but small so this is a fairly competitive market.

In that case -- in a unique symmetric pure strategy equilibrium -- allocations are first best. There's no distortion, although firms are charging different markups, and there are surprise penalty fees. And firms endogenously choose not to disclose at the point of sale whether or not they apply, and prices could include where the best marginal fee is zero, but penalty fees are high so they expected marginal prices are equal to marginal costs.

So what this means is if we impose price posting regulation, that nice efficiency that allows the pricing goes away so it's counterproductive. Having price posting regulation would lower welfare. It would hurt firms. It would hurt some consumers but benefit others, and I'm going to skip the intuition for that unfortunately, but my interpretation here is that the bill shock regulation should be applied with caution. You really need to be worried if you think people have correct beliefs, and it is a fairly competitive market.

I'm going to skip this. So in the third model, I assume that people underestimate their demand for the service. If people are unbiased, the thing about a monopoly now is they're all the same so there's no price discrimination. I would fit marginal pricing into the marginal cost so they get total surplus.

The blue bar would be at first best, the dashed line, firms have captured all for a fixed fee, consumers would get none. If consumers underestimate their demand, we can't charge them their full surplus upfront for a fixed fee. They don't anticipate it, so we have to charge it, capture it through high marginal fees. Allocations are distorted. Total surplus goes down, and firms can't capture all of it. Consumers get some.

With inattention, in addition to this bias and beliefs, it's ambiguous whether total surplus goes down or up, but consumers no longer need to get at least their outside option. They can be exploited in the sense that their utility can be negative relative to their outside option. Firms can capture more than the entire surplus, so it's ambiguous whether the regulation in this case would be good or bad for total welfare, but it would definitely protect consumers from this type of exploitation.

So one of the things that's important to know here is if we have a good with no social value or if people underestimate its value, it's not going to be sold so everybody gets zero, so they're unbiased. But if they're biased, the underestimated value is attentive, that remains true, but if they're also inattentive, you could end up having a business that starts up and starts selling a product with negative social value just to earn money on penalty fees.

So that if the social surplus is negative, consumers do very badly and the firm actually makes money. I think that's interesting to note because when this new federal regulation came in and required banks to ask consumers, "Do you really want overdraft protection before they started giving them services and charging the fee," Bank of America, who had been earning about $2-billion a year from these fees, decided to end the service rather than asking consumers if they would like it. So this is possibly one explanation.

So to conclude, when consumers are inattentive, if you think everybody is the same, there's no price discrimination going on and if people have correct beliefs, then this is not a non issue. If there is scope for price discrimination because people are heterogeneous in fairly competitive markets, inattention combined with penalty fees can be socially valuable, and disclosure regulation could be counterproductive.

When people have biased beliefs, it's ambiguous whether the regulation would be good or bad for total welfare, but it might be that the larger factor is that the regulation protects people from exploitation. I think these two results fit nicely to cell phone pricing and overdraft fees.

DR. LAIBSON: Thank you. We now have Ginger Jin to discuss the paper from the University of

Maryland.

DR. JIN: I really appreciate the opportunity to read and discuss this interesting paper. The question Mike asked is very simple: What happens if consumers do not pay attention? This is a very timely question for many consumer protection policies, and Mike tries to be very ambitious and comprehensive in this paper.

It was like after I turned the page, I would say to myself, "Oh, I bet he doesn't mention this," but then he mentioned it in the next page. Not only does he try to capture attentive and inattentive consumers, he's also considering whether the consumers have a correct or biased belief about future demand and whether the future consumers are homogenous or heterogeneous with competition on the supply side, a monopoly to competition and different types of consumer protection policies that we could imply in this market.

So when I read the introduction, I thought there were some results that are pretty intuitive and some results that are surprising. The not surprising ones are that firms for sure are going to exploit consumer inattention and underestimation of demand. Consumers would exercise price discrimination whenever it's possible.

What's surprising is the equivalence results when consumers are homogenous and have the right belief, which was very surprising to me when I first read it. Of course after the model, I find it intuitive that firms can replace the penalty fee with other charges and don't change anything.

A more surprising result is in the fairly competitive market with heterogeneous but unbiased consumers actually allowing penalty fees would be good for the whole society. My take is that that's because with the penalty fees, it gives more room for the firm to price discriminate in a less distorting way, but when we restrict the potential pools the firms can use for price discrimination, then it could generate more distortion in the process of price discrimination.

The last message I'm taking from the paper is that the transparency regulation could be more important in terms of redistribution of surplus between the firms and the consumers, rather than enhancing the total welfare.

So, I have some comments. The first is exactly how should we think about consumer inattention? Mike has been very clear about what he is considering. He's considering consumers who do not pay attention to their past usage. However, consumers are fully aware of their future inattention, so they take that when they sign the contract, so they have an expectation of the future penalty fees and so forth. And they also have a belief on future demand on whether that belief may be correct or biased.

I can think of situations that are probably more general than this considered set of consumer inattention. For example, when they sign a contract, consumers may not pay attention to all the contract terms, including the fine print and so forth, and even if they pay attention, some contract terms could be waived or hidden.

"We reserve the right to change the price in the future." Well, what does that mean? And consumers may not realize the risk they're exposing themselves to when they sign a lock in contract. People would have a belief on the future demand if it appears that we can impose a probability of distribution on that, but if we walk away from that assumption, it's possible that the consumers don't know what kind of distribution they should put on that future demand.

Another comment is on the consumer protection policy. My take of the whole paper seems like the need for consumer protection policy is driven by consumer underestimation of the demand.

Well, supposing that is true, does that mean that we should actually target consumer education instead of firm regulation if you think of sort of the reason driving for this kind of inefficiency? If consumers do not pay attention, maybe we should educate them? While 20 percent of people who use our service actually end up paying penalty fees somewhere down the road, that will be an affirmative message to me so that could be alternative policies, not just banning firms from imposing penalty fees and requiring firms to post the price.

A deeper question I would like to ask is: Why do consumers underestimate demand? Of course there are allocations in the future I cannot foresee right now, so that could be how I could underestimate it or overestimate it. On average I may be right, but another reason could be I'm not familiar with the future service. Like when I purchaser cell phone for my husband, he was 100 percent sure that he didn't need a phone at all, but of course one month later, he could not live without a phone.

So you just don't know what you're buying yourself into, and then your estimation about the future demand could be way off the mark. And so that's kind of saying that the demand under that estimation is something that's maybe not completely exogenous.

The third one is if I'm thinking the consumer protection policy would reduce the penalty fee and reduce the surprises down the road, it's encouraging consumers to be attentive. Would that have long run consequences, like we'll have more people talking on the phone while driving and we'll have more people overreacting to their bank account and -- probably several years later -- file personal bankruptcy because they didn't take care of their finance when they should have taken care of it?

So those are some things that I know are not in Mike's model, but when we consider the consumer protection policy, maybe that's something we should take into account.

I can think of several directions to extend this already comprehensive theory. One is attentive versus inattentive is an endogenous choice, but I know I may not have time to pay attention to my electricity bill or cell phone bill, but can I choose to opt-in to the price posting regime before I sign a contract so that I force myself to pay attention to it?

If the underestimation of demand was driven by the addictive nature of the service, then the penalty fee might in some ironic way limit that addiction.

Some factors not considered in the paper might be important, like risk aversion, how long I'm locking myself into the contract, or what the switching costs are. If I dislike the surprise penalty fee can I switch out with a reasonable cost? All those help interact with inattention.

The last one I think Mike mentioned a little bit in the paper is that the penalty fee and restriction definitely are related to the design of different usage plans. So is it because there are so much penalty fees for calling over my limit, that it will force me into signing up for a very generous usage plan, which ends up charging me a lot of dollars?

Overall it's an interesting paper. I really enjoyed it. I find the results very stimulating, and I would encourage everyone to read the paper themselves.

Thank you.

(Applause.)

DR. LAIBSON: Thanks very much. We're going to move on to the next paper and reserve questions to the end.

Now Jonathan Zinman from Dartmouth is going to tell us more about inattention.

DR. ZINMAN: Great. So this is joint work with my frequent coauthor, Victor Stango at Davis. This is not the first time I've presented a Stango and Zinman paper at the FTC, and I hope it's not the last. It's very good to be back.

So what we're working on in this paper is we're interested in limited attention and its dynamics, so located and varying consumer attention with regard to the payment of bank overdraft fees.

And so just to fix ideas -- or actually in this case, keep ideas rather vague -- our working definition for limited attention for the purposes of today will just be that people only imperfectly integrate information on their choice sets into their decision making. I'll talk a little bit later about how we might be able to tie some of our results to different types of theory models, but for now let's be agnostic.

So what we do in this paper is we use subtle variation in survey content, in the questions that people that happen to be taking surveys are asked as potential shocks to attention with regard to payment or incurring bank overdraft fees. So we have this panel of transaction level data on consumers in our panel, for reasons that I'll describe, of frequently offered service, the topics of which I have not announced in advance, and so these surveys have questions that mention overdrafts.

Just to preview what we found, our attention delves into whether attention has an effect on overdraft fees, and we find a large reduction in overdraft fee payment following what are often relatively subtle attention shocks in the form of one or two questions on a longer survey.

We find evidence of economically important dynamics. Attention seems to accumulate so if you're exposed to repeated shocks, if you take multiple surveys over time that mention bank overdraft fees, your baseline level of overdrafting drops, but both of these effects -- both the immediate effect and the stock effect -- depreciate over time.

These effects are variable. They're heterogeneous, and they are largest for some groups that are viewed as being "particularly vulnerable" by some policy makers and consumer advocates. Affects are largest in low self assessed sophistication folks. We do not find differences by high versus low income.

So what we're not doing at this point is trying to say anything about the welfare implications of our results, so I just want to state that upfront to hopefully put your mind at ease about the types of claims that we're making or -- more to the point -- we're not making.

It may be the responses we find and the underlying limited attention we think that are total effects in our response. The responses to these surveys, we think, indicate the underlying limited attention that may or may not be suboptimal. If I have time, I'll talk a bit more about that at the end.

Some quick motivation. Michael has already given us some, and this audience is probably relatively well schooled on the economic importance of overdraft fees, but nonetheless these are a major expense for U.S. consumers, so there have been some nice and provocatively put together accurate statistics showing that U.S. consumers in recent years have been steadily charged more on overdraft fees than they have for various types of fresh produce or even on large appliances, and we certainly wouldn't want that to be the case.

Overdraft and overdrafts and overdraft fees mean different things to different banks and under different contracts, but what you should have in mind for today is an overdraft is basically a transaction that if it is settled by the bank would result in a negative balance in the checking account. So it's a transaction that if it's settled produces borrowing against future checking account balances, and the typical fees in recent years for this sort of transaction have been upwards of $20 per transaction, per loan, and this is regardless of the amount of loan.

One reason we got interested in the question of whether limited attention might play a role in this -- even before we were lucky enough to have Michael start working on this - is, I think, that this seems like a market that is very much an architectural attribute and firms go out and advertise "free checking" and distract attention from or do not mention the fact that free checking is not free if you incur a $35 overdraft.

The first thing we did when we got this data that we're using in this paper is we put together some descriptive stats for our papers and proceedings paper. And we found some evidence there suggesting that many overdraft fees are easily avoidable in the sense that many people pay fees at a point in time where they have much cheaper and readily available sources of liquidity.

So they're using their debit card to pay for a transaction at point of sale, and they're incurring a $39 overdraft fee when they could have paid for that same transaction with a credit card that we observe in this data as having available liquidity. They could have paid 30 cents to charge that transaction instead of 39 bucks.

There's some evidence in our data here of limited attention to balances of the type that I think motivates Michael. If you ask people why they overdraft, 60 percent of folks say, "Oh, I thought there was enough money in my account."

We'll go on to the bank side. Overdraft pricing and overdraft fees have been something that banks actually consistently did well throughout the 2000s, so it's been a major profit center for them by various metrics.

I think this is the final piece of motivation: There's not a ton of evidence on the sort of empirical determinants of supply and demand in terms of nonlinear stake attention contracting, particularly in household finance.

One of the things we hope to do in future versions of this paper is speak more directly to Gabaix type models and to Grubb type models as you will see. I think our evidence model ultimately will have something to say about the types of heterogeneity or homogeneity in Michael's model, about the types of biased or unbiased beliefs that's critical to Michael's model, so that's the direction we hope to push.

On the limited attention to memory side, there's a host of different theory models out there. The amount of paper that tracks that side of the literature, a different paper as well, our comparative advantages here are not just providing empirical evidence but having evidence on the dynamics. And the very last substantive slide I'll show you today hopefully will be on the mechanisms about how people go about reoptimizing, how people go about implementing these reductions in fee payments.

And the fact that these surveys we use as attention shocks have effects also relates to a couple of literatures on priming and on how surveys change behavior. The data we have is 36 months of panel data from checking account statements, and also from credit card statements although we're not using the credit card data selection in this paper yet.

This data is put together by a market research firm, Foster, which is well known to economists. They actually broke off this piece of the business. It's now owned by Light Speed Market Research so we have over 7,000 panelists with active checking accounts that we used in this paper, and over a hundred thousand panelists months worth of data.

So panelists entered the data typically after having some other relationship with this market research firm. The firm goes to them and says, "Hey, we will pay you let's say a remarkably low amount, 20 or 25 bucks if you will sign over access to your online account statements." People do so again presumably in large part because they've interacted with the market research firm before and trust them.

Once they sign over access, the market research firm goes through there and scrapes data from their account everyday, so that that's how we get the account and transactional level data. Again because this is a market research data firm, when people agree, the market research firm has them take an online registration survey.

We get some demographics and our self-assessed measure of financial sophistication from that survey, and then periodically the market research firm offers surveys on a roughly quarterly basis, although it's not quite that predictable. These are all market research surveys about people's financial relationships and vendors and satisfaction levels and what they think of new products, both hypothetical and actual.

These surveys are all online. They are not lengthy in the sense that the market research firm makes a big effort to keep them to about 15 minutes in length, but because they are online surveys and a lot of the questions are simple, they can contain as many as a hundred questions or a couple hundred questions.

So then we observe our panelists' full survey taking history, not just within the three-year period over which we have transactional data, but also back a couple years beforehand. Important for our empirical strategy is these are not pre announced, so people presumably have some idea that this is going to have something to do with household finance, but people get sent Email invitations to participate in these quarterly surveys, and all it says is, "Click through to take a survey." It doesn't say what the survey is going to be about.

So let me say a few words about external validity. Clearly this is a whacky group of folks who have lower reservation prices for signing over access to their sensitive financial information. It's people who took the detailed household surveys that we use the SCF or the SE analogs in the EU, are probably somewhat similar along unobservables if you think about it.

I mean, these are people who have again low reservation prices for revealing very sensitive financial information to surveyors who they've never met before. At least our guys have interacted with this market research firm, so a whacky group of folks on some difficult to observe or unobservable dimensions focusing on observables. Because these folks are online perhaps they tend to be younger, more educated, higher income than the U.S. average, a bit more creditworthy condition.

An interesting thing to note about the survey is that by any of the metrics we usually use to proxy for financial sophistication are samples, our sample is relatively sophisticated. But I don't know of a clear prediction on whether more or less sophisticated folks would respond more or less to the types of subtle attention shocks that we're interested in here, so just something to keep in mind.

So there are 21 of these roughly quarterly surveys. Six of the 21 have typically one or two questions that mention bank overdraft fees. We think of these surveys as potential shocks to the attention or salience of overdraft fees. Everything we do is within the panelists and conditional on selection into survey taking or not survey taking generally.

Lots of people on this panel take surveys.

Lots of people take lots of surveys. Five of the six surveys basically mention overdrafts, again one or two questions. I'd give you examples, but I have to speed up here to get to the results.

When we say they mention overdrafts, what we have in mind is there's no mention of prices or outside options. There's plausibly no information being provided on what is in people's choice sets or what might be in people's choice sets. And there's this sixth survey, which really does beat people over the head with lots of questions on overdrafts and plausibly does provide some information.

So again I'll skirt over this in the interest of time, but basically we're going to be looking within panelists across months and looking at how overdraft fee payment responds in the month that you take a survey and also as your stock of overdraft related surveys builds up over time. And this is all conditional on whether and when you are taking our other surveys that do not mention overdrafts.

So what we need for our identifying assumption is that conditional on survey taking, there are no differential, unobserved, secular dynamics in overdraft fee payment across those who take relevant surveys and any other survey. And so the dynamics that would have to be present to confound us would have to be high frequency given the timing of the surveys and the nature of our findings. Again the survey topics are not announced ahead of time.

So we find a big immediate reduction. This is the same month effect, larger point estimates for the least educated guys and significantly larger reductions for the low self assessed financial literacy guys. We find that overdraft fee payment falls as the stock of taking these overdraft mentioning surveys builds up.

We can go and see because there are lots of different types of surveys and lots of different survey content, and we don't find that overdraft fee payment responds in the wake of people taking surveys that mention gift cards or auto loans and so on and so forth.

We do have some findings on related content, and it is the case that overdraft fee payment does change following taking surveys that mention other bank fees. The flip holds as well, so fee payment generally falls following surveys that mention bank overdraft fees, and so we think this tells us something about the cognitive process associations in salience.

So in beginning to unpack these cognitive mechanics, it's important to keep in mind these surveys can be affecting behaviors in several different ways. One is what we most have in mind, which is an attention shock or a reminder of the price schedule you face, but another particular point is, given the often inadequate upfront market disclosure selected, that there is some information being provided. This is particularly plausible in one survey that has a bunch of questions on overdrafts and both reminders and information could be at work here.

So what we find here is that that more intensive survey does have an incremental effect. And so does that mean the more intense survey is a more powerful attention getting treatment, or is this an incremental effect of information? We haven't figured out a way to unpack that yet.

We talked about how effects depreciate over time. We only find effects on the extensive marginal fee payment which is consistent with attention being limited in a discrete way, and I already talked about attention by association and those roles.

So some new results that I would spring on Karen, -- Sorry, Karen, but we got the versions of these in late last night - relate to how people pull off these fee reductions. We're finding that people are in fact spending less out of their checking account so you might imagine that people just transfer money in or transfer less money in, but what seems to be going on is people are actually spending less.

This includes people who never overdraft throughout the sample, which is interesting to think about. Spending falls particularly as balances get low, especially for the never overdrafting guys and most of all for frequently overdrafting guys, but spending falls more globally for everyone. So this raises, I think, some interesting questions about what the marginal transaction is as well as who the marginal inattentive consumer is.

One thing we'll be looking at going forward is do people hold their spending constant by reallocating efficiently to their credit card accounts? We'll be able to measure that eventually. We don't find any effect on balances, which again suggests this is not about people moving money in or out of the account.

So just to sum up, our results suggest that consumer attention to this kind of state contingent penalty pricing is limited, discrete, dynamic, malleable, heterogeneous in its malleability and associative. And since I'm out of time, I'll pass on the speculation of implications for disclosure policy since Michael and Ginger have already done that ably.

Thanks.

(Applause.)

DR. LAIBSON: Thank you. Karen Pence will be the discussant. Karen is with the Federal Reserve Board.

DR. PENCE: Let me point you to the most important sentence in this whole presentation -- the one thing I do want you to take away -- which is the disclaimer. I'm here in my personal capacity. I'm not representing the Federal Reserve or its staff, and my remarks do not represent those of the Federal Reserve.

So, this was a fun paper to read. It's a fascinating paper. It's a fascinating market. It's one I had not thought very hard about before. I confess I still don't know how my own financial institution handles overdrafts.

Overdraft fees are a huge source of revenue for banks, so this similar to what John just showed you displayed a little differently. In 2008, overdraft fees were $36-billion worth of revenue for banks, so that was over half of the revenue they got from checking accounts, and it is fees that are paid in a very disproportionate way. This is from John's paper; it may look familiar to him. This is a distribution of the share of accounts that have overdrafts and for how many months they have overdrafts.

So it's a huge amount of people, over 50 percent, that never pay an overdraft fee, and then there's an enormous tail of some people who pay quite a lot of fees. I think John had a statistic that there are people that pay up to $250 in fees in overdrafts, which is enormous. So understandably, this has been the focus of a lot of public policy concern.

To briefly reprise the findings: bank customers that are reminded, and the way they're minded is by participating in a survey about the overdraft features of their account, are less likely to subsequently incur overdraft fees, and it's a larger effect in financially vulnerable groups.

The results are still preliminary. One thing I think that's a little bit unfortunate is that it would be nice if the authors had more variation in when the overdraft surveys were asked. They're very concentrated on this short period of August to November 2006, even though their transaction data spans a much longer period of time. The authors do everything they can about that. They have month year fixed effects. They have person fixed effects. Nonetheless, I think it would be a little more comforted if they were not so concentrated on one point in time, given how much has been going on in the economy over the past few years.

Nonetheless, if you can't see my last point here, the finding is consistent across samples and specifications, and they do quite a lot of robustness tests and falsifiable tests so it's a fairly convincing result. And as has been previewed, the Federal Reserve has done all researchers a tremendous favor, including John, in bringing about these revisions to Regulation E.

So this is a wonderful natural experiment in changes in overdrafts and opt-in defaults and messaging. There are two key dates. One is July 1. The second is October 15, but as of August 15, banks could no longer provide standard overdraft services for everyday debit card or ATM transaction, unless the customer explicitly opts in.

Now interestingly it's not everything, so if you write a check and that bounces, you overdraw your account, the bank can still provide you overdraft services. If I have a recurring payment set up, that was not covered. It's just the scenario people talked about. You go to Starbucks, and you slide your debit card through, and you didn't know your coffee was going to cost you an additional $35. So those are the kind of transactions envisioned. That's what everyday means, which is now a federal regulatory term, I think, in this context.

So, banks provided extensive messaging to their customers about this opt-in choice, and I'll be showing you a little bit about that messaging later, but again there's a lot of interesting stuff here in terms of communication. One nice thing about the data that John and Victor have is that they have very, very detailed information on the transaction, so in theory, they could look separately at check transactions, everyday debit card transactions, that kind of thing.

There's this really neat statistical relationship in the data. You take a survey. You're less likely to incur an overdraft. So for policy makers the crucial question is going to be: What is that mechanism? It's a neat result, but what should I draw from it?

John highlighted two possibilities. I'm going to talk a little bit about one of them, which is information, and you can think of information in a couple different spectrums. You can literally be not aware that you're going to be charged a fee if you overdraw your account or you can be unaware of certain aspects of the policy.

Then they're going to test this in this one survey module, which is much longer, much more extensive. It has all this fee information. It has 12 questions on overdrafts, but the behavior of people who took that survey, was not affected anymore than people that took saw a single question, "Do you have overdraft," so that makes it hard to think just based on this limited information that it's information per se that is the driving factor.

I'm going to show you a slide that I see as consistent actually with his view that information can't be the main factor. This is a slide from a consumer research group called Mintel, which is fairly similar to the one that John is using in his paper. They also do Internet surveys, so this is a survey in June of 2010 of a thousand adults over the Internet.

They were asked two questions: First, have you overdrawn your account in the past six months? And secondly, are you planning to opt-in, so have you already opted in? Are you not planning to opt-in?

What I find interesting is I think incurring one of these fees is probably the best form of education you can have. There's nothing like paying $35 to make it very salient to you that overdrawing your account is very expensive. No amount of reading disclosure is going to convey that message that vividly.

If you look here, this is all customers, and these are the ones who overdrew their accounts. Of the ones who overdrew their accounts, they're the ones that are choosing to opt-in. I don't know if they know what they're doing, but they're very clear that overdrawing brings a penalty. There's not a basic educational problem or informational problem, and they want the service.

So this suggests to me that some of the things John was talking about, about salience and about reminders, giving people the information so they can make a decision in advance as opposed to having it sprung at them at the register, is actually the way to think about it. It is the salience. It is the reminder feature that may be the round in their data.

Just to finish, while we're on the subject of salience and reminders, we often get from Mintel copies of all the mailings that banks sent to their customers urging them to opt-in, and the most frequent message is actually of embarrassment. I'll read you three quotes. "Let us continue to save you the embarrassment of having your purchases declined and the hassle of not being able to get cash in an emergency."

"The benefits of overdraft protection: It's convenient. It saves embarrassment. It provides a safety net. It's good to have in an emergency."

Finally my favorite: "Our intention has always been to save the embarrassment and inconvenience of a declined transaction."

So my point here is that salience reminders are something in this paper promoted as a good thing so they can help consumers make the right decision. I think it's important to remember there are tools that can be used to influence decisions in a whole lot of ways, and this is a pretty stark reminder of that.

Thanks.

(Applause.)

DR. LAIBSON: Thanks. The third paper now, Nicola Lacetera will present about limited attention in the new car market.

DR. LACETERA: Thank you very much. It's so good to be the third presenter on consumers because I don't have to spend too much time saying consumers are inattentive and that that seems to matter in important markets.

What we do in this paper with my coauthors, Devin Pope and Justin Sydnor, is looking at a particular market, the used car market, which as you know is a very large market with large stakes. It's reasonably competitive, and yet the particular type of inattention we will look at is not washed away by the used car activities of this market.

But the type of inattention we look at is inattention to the exact mileage of a car. Mileage or odometer reading is, on the one hand, a very important determinant of the price of a car, and on the other hand, it is somewhat different from a good part of the literature on the effect of the inattention.

There isn't this obfuscated or shrouded component to it, so it's a fully visible thing, so we can look at the odometer in the particular setting which is a wholesale used car option. It's going to be very clearly displayed on a monitor, so anybody can see the exact mileage of the car, and yet it seems to be the case that there is not full attention to the mileage. And in particular we will look for this left digit bias whereby we pay full attention to the left most digit of the number and less so to the other digits.

Just to organize our empirical analysis, we use a framework which has been developed by Stefano Dellavigna based on others' works. The idea is we have two sets of characteristics affecting the value, the perceived value of that good. One is a fully visible characteristic, V, and the other is an opaque component, so we cannot see it immediately, and we don't pay full attention to it, so the attention that we pay is determined by the factor one minus theta.

So theta in this is our measure of inattention. How do we translate that in our survey? We look for this left digit bias. Essentially think of a number like 49,900. Think of a mileage of that type. The way we think about how consumers look at this number is essentially by discounting the right most digit in a constant way, one minus theta, even in a progressive way.

So the more we move to the right, the less attention we pay, and there is some psychology research in this, in the form of recall surveys that confirm that this seems to be the case.

How does this translate into a valuation schedule? So, let's think about the car market as an example. You think of the perceived value of the car being determined negatively by the mileage, and if we believe in that inattention where we assume in terms of valuation this discontinuity in the simplest case at each and every 10,000 mile mark, it is proportional to the level of inattention, of course, and to the depreciation rates of the car.

Essentially we perceive the mileage to be lower than the actual one. The valuation is always higher than the actual one, which is represented here by the dotted line, and it coincides with the actual one, at the exact 10,000 mile mark.

How does that translate into auto sales? So we will look at wholesale car auctions where sellers -- they can be dealers, they can be car rental companies, and so on -- bring their cars to the auction. On the demand side, we have used car dealers who will bid for these cars in an ascending first price auction scheme, and then they bring their cars to their lots and sell them to the final consumer.

So in the simplest kind of model, there is this representative agent in the competitive market, and then we will have the price at the auction coincide with the prices on the lot because of the zero profit condition. And the price will reflect the perceived valuation as we show sort of in this graph. We can enrich the model by considering the heterogeneity of consumers by valuation or by the level of inattention, but the results in terms of the gaps we will have stayed the same.

Empirically, we looked at the major wholesale used car auction companies in the U.S., which gave us access to their data. I was almost sure we had the largest data set in the room, and I think they beat me like five to one so that's unfortunate, but 27 million observations seems to be a pretty large data set, so we were excited about it.

Just very briefly, what we got was the information about the car sale, the price at which it sells, the exact mileage, which again is displayed when a car is auctioned, and those are also available in printouts, and we know, importantly, a lot of characteristics of the car: the make, the model, the body style, the production year. We know the year of the auction, so when the car was transacted and also the location of the auction. We also know the precise ID, essentially the identity in a sense of the buyers and the sellers, so we will use all of this information in our analysis.

So let me start by giving you a sense of what the whole data looks like. So what happens if we plot on an X axis mileage and on the Y axis average price? We see that something is going on, right?

We see that there is a decline, of course, and we see discontinuities pretty much in each and at every 10,000 mile mark. If we go and look very closely, we see that also each point represented a 500 mile bin. In a sense there seems to be a little bit of inattention also in the sort of four digits, if you want, the small ones.

So there seems to be something, and of course we can from the raw data interpret this exact mileage so we can actually give some numbers. For example, as we cross 80,000 miles, we have a difference of about $200 while as we stay very close to 80,000 miles, but a little bit under 70,000, it's just $10 for each hundred miles of the older cars.

So can we assign this to inattention? Is it the right size? Maybe there is some selection going on? So, for example, the seller might anticipate that this will happen, so they bring different types of cars on one side or the other of the 10,000 mile mark.

To the extent that this car differs on features that affect the price, we can actually have some buyers estimate. And something like that might actually be going on if you look at the volume patterns. So this is the number of cars brought to the auction by mileage, and you see some weird things going on.

I'll explain in a moment what is going on around the 30,000 mile, but in general, we see some peaks right before each 10,000 mile mark suggesting, in a sense, that maybe there is some timing, so to speak, in which cars are brought to the auction.

So one way to sort out this selection effect is essentially to look at receipts, so we had a lot of information about the cars. We can construct this very detailed fixed effect for each car, make, model, model year, body style, up to the auction location, auction year and even the ID of a given seller. We can look at the regression of the price on these fixed effects, and this discontinuity stays there.

By the way, doing this kind of analysis is the same as running a regression discontinuity kind of analysis in which essentially those discontinuities are equivalent to the estimate of these dummies on these 10,000 mile mark changes. Again we see that also for the thousand miles; we see this kind of pair. It's very difficult to see -- it's small of course -- but it's there.

We also decided to go after the issue of selection by looking at the two different classes of sellers we have in this data. So, on the one hand sellers are new car dealers who are trading in and they decide to bring it to the auction. On the other hand, we have Hertz, Enterprise or financial companies getting rid of their leased cars, so for the dealer sellers, we might be concerned about selection and what kind of cars they bring to the auction as opposed to those that they keep on their lot.

They may have a higher reservation price, and, for example, they might be much less likely to sell a car. It's actually about 70 percent of the cars sell as opposed to the other type of sellers, fleet and lease as the company has classified them. They actually sell almost all of their cars just to get rid of them.

So selection issues seem to be less of a problem there, and in fact we see the spikes in say bringing a car to an auction for the fleet/leases cars around 36,000 and 48,000 miles. That depends on the structure of the lease, four-year 48,000, or a three-year 36,000 miles, but these peaks are typical of the fleet/lease cars, those peaks in volume right before the 10,000 mark, so there might be some selection effect where it's important to control for the fixed effect.

Once we do that and we see that the patterns are very similar across the two types of sellers, so these are the same procedural graphs, but separating the two types of sellers. It also is important if there may be other alternative explanations that we need to take care of.

So first of all, one could say, "Okay, you convinced me that all the unobservable characteristics we need to account for this selection; how about unobservables?" We believe that in this setting, prices are determined by the characteristics that we can actually see in our data, so it might not be too much of an issue.

We still need to be concerned about those peaks right before the marks, but the whole schedule of price actually shifts down, not just around the marks so that those will make us less concerned.

Of course a major issue is warranties. Maybe people just discount the warranty structures as an instructional characteristic. Fair enough, but we don't observe what it is for all of the 10,000 mile mark. We also analyze separately makes from which we have more detailed information on the warranty schedule, and we can rule out essentially that explanation as well.

Maybe people cheat, and they tamper with the odometers or they bring it back a little bit. We don't have empirical evidence of that, of course. This is something they don't tell us about, but this will bias our estimates downward. These people anticipate there's some probability that this might be going on.

What about the differences across time? We do see some heterogeneities across cars, right, so for some cars, the discontinuity is higher than in others. Actually the simple model presented gives an explanation for the idea that the higher the depreciation, the higher the discontinuity should be. And that's what we find in the data.

So then we wondered: Well, maybe people just look at the books like Kelly Blue Book and Edmunds, and those books are discontinuities, so they just follow what they see there. Again for Edmunds, they actually use a new procedure where there is no discontinuity. With Kelly Blue Book, you will see if you plot data there are some discontinuities, but it's not systematic, not at the 10,000 miles mark.

So again even the sort of institutional explanations don't explain this. Since now I'm employed by a Canadian institution and my discussant is Canadian, I thought that I should say something about the Canadian data which we rely on just for a few auctions. As it turns out for the Canadian auction, the discontinuities are at the 10,000 kilometers mark, and we run it as a perceived 10,000 mile mark as well.

We don't see any discontinuity in the 10,000 mile mark, and this is consistent with this left digit bias.

Now, it's kind of interesting to understand: Who is inattentive? We have different buyers at the auction, and the results we find in particular are consistent observationally with two cases. One is in which the buyers at the auction are fully savvy and they just anticipate the final consumers not being attentive, and another is where the buyers at the auction and the final consumers share the same bias, and that's where we cannot tell them apart.

We do not have conclusive evidence because of these observation equivalents, but what we can do is to look at more experienced buyers. Those that come more often to the auction are more likely to buy before the threshold, so they do not perceive it as overpriced, and they can anticipate they can sell the car for more.

We also observe, if we look at the data, that the drop begins actually before the 10,000 mile mark, and you have to drive your car back to the lot, and there are going to be some test drives. You want to be sure by the time the car is sold that it is still below a 10,000 mile mark. And we actually called and talked to some of these dealers directly and they said, "Yeah, actually we make sure that when we test drive cars, they don't go over the 10,000 mile mark."

So we conclude that most of these buyers should be on the side of the final consumers who are then the ones that bear the mispricing implied in these.

Finally we would like to give an estimate to the amount of the inattention or the degree of attention, the theta parameter in that model. We go after that in different ways, so as you can see, there are very simple linear ways in which mileage entered the equation.

We can estimate the theta by looking at the gaps and estimating the depreciation and back out the theta, which is around 30 percent. Similar results use a nonlinear specification with a flexible polynomial in miles rather than the heterogeneity in the price discontinuity.

If you essentially regress the gaps you observe on the side of the discontinuity, the coefficient that we estimate should be our theta, and again we find it to be in around .3, which implies that 30 percent of the price decrease can be explained by this gap at the 10,000 mile mark.

So to conclude, we find evidence of these buyers being present in the market, being in the order of discontinuities between 150 and 200 dollars. This is mispricing if you integrate 2 to 3 billion dollars and if you believe that it is the final consumers and these buyers, they're the ones from which the welfare is effected. We also see an effect on supply decisions because we see these patterns, so it seems that these buyers affect different aspects of these markets.

One could ask: Is it a rational thing to do, not to pay attention? There is some cost and people factor in these costs. In this case, in this setting, it's easy to look at the odometer. You buy a car with 9,999 miles, and you know that in a matter of few days, you will cross the 10,000 mark. It's also true that we wouldn't expect the discontinuity to be bigger for cars with higher depreciation because the cost of the possessing that information shouldn't depend on both, but I think the evidence is more toward an irrational type of inattention rather than a rationale one, but we don't have fully definitive answers.

Of course we speculated that there are many other settings where it happens. The numeric measure is important in decision making. Think of GPA for hiring or SATs for admission decisions. Think of accounting measures for financial evaluations and so on or medical measures, pressures, weight of babies, newborns and so on.

That's where understanding better this type of left digit buyer and the effect on market and decisions would be totally relevant in future research.

Thank you very much.

(Applause.)

DR. LAIBSON: Kory Kroft will be our discussant from Yale.

DR. KROFT: Thanks for the opportunity to discuss this paper. There are a few things I would like to say at the outset before getting into some detailed problems.

The first is that there's growing literature on the importance of inattention, and the majority of studies in this area come typically from lab experience, lab experiments, and this paper is different in the sense that it focuses on a naturally incurring market and uses observational data.

So, one of the contributions is estimating the importance of inattention in an equilibrium setting, focusing on the steady state, which I think is interesting.

The second departure from the previous literature is that typically other papers have focused on situations where information is shrouded and typically there's been some experimental manipulation, and in this context, information is not shrouded, in the language of Zinman and Laibson. So, it's more impressive if they find that inattention is important because information is available at no cost.

Apart from the fact that the evidence is extremely credible, the last thing I wanted to say at the outset, or at least one thing I liked about the paper, is they had a really nice simple framework and intuitive behavioral model for interpreting their evidence, which delivered a parameter with a nice interpretation, and others have estimated it so that they can see how inattention in their context relates to inattention in other contexts.

A few things I would like to discuss and I won't grapple too much with the empirical evidence, but I wanted to talk about how we interpret some of the estimates, and I wanted to say something about welfare. Nicola talked little about this volume response, so I won't talk about that, and finally we'll conclude with the empirical test of final consumers versus used car dealers.

So the claim of the paper is that theta, that they interpret this by saying 30 percent of the depreciation that a car experiences due to mileage increases occurs discontinuously at 10,000 mile thresholds, and I don't want to push too hard on this, but there are sort of several caveats that I wanted to mention.

The first, and I could be off on this, but the way that they're getting this estimate is basically to compare the discontinuity to the rate of depreciation, and I wonder whether one wants to estimate the rate of depreciation using raw prices or residuals. And I think the way they're getting the depreciation rate is basically they're holding the age of the car fixed, so I wonder if you want to use an unadjusted or adjusted price. So, I wasn't sure about that.

Then the other comments with respect to the interpretation of theta is it sounds like 30 percent relies to a large extent on the model, so the way I think about this decision is going out and buying a car and having a decision where you think you're going to drive that car for a couple years or 30 or 40,000 miles, and then at some point you think that you're going to sell the car.

And it's at that time where the margin of whether to sell it at 49 versus 50,000 comes into play, and if you think about that sort of a model, then inattention and that kind of a decision problem seems more like sort of a local phenomena in explaining depreciation.

In the estimates for accounting depreciation, I think you might want to consider outside the model to what extent it accounts for the depreciation of the car and also the potential welfare cost.

So Nicola mentioned that not only do we observe a difference in price for cars below mileage thresholds and cars at the mileage threshold, but also observed sellers responding to the threshold in the sense that sellers were more likely to bring their cars to the auction right below the threshold.

And if you think of that as a supply response, one question is whether that effects the equilibrium price per car, so if sellers are responding endogenously to a higher price increase, then one would think that that has some effect on the price of the car. And I know they came at from the standpoint of focusing on whether there's a selection effect which I think is important, and I'm, kind of convinced that there isn't a selection effect going on.

I wonder whether the price difference might even be greater once you account for the downward pressure due to the supply response.

Finally, it might be interesting just to study the volume discontinuity of dealers as an outcome rather than something to look at in terms of a selection problem. You might think that's a function of the parameters of the model, so it might be interesting to think about how to model the volume discounts or in addition to the price discounts really.

So in the interest of time, I'm going to skip over some of this. I'll just talk about some related research ideas that you might be able to explore in the data.

One of the nice advantages of this data set is they have data from the U.S. and Canada, so one thing that Canadians talk about is if you take two cars that are identical in terms of the model and make and mileage, and one car is expressed in miles and the other one is expressed in kilometers, do you observe a different price just given that the one car is presented in kilometers?

So I don't know if you have enough power in your data to test this, but it could be interesting to see conditional on what make and model and mileage, whether just changing the metric that the mileage is presented in has an independent effect on price. And I think it fits in with your attention framework.

And then another phenomena that people talk about in car auctions is that price seems to be based on the age of the car in terms of calendar years rather than months. You could imagine comparing two cars, one purchased right at the end of the year and the other purchased right at the beginning of the year -- two cars that are kind of observationally the same -- and see if you find any difference in price.

Again you have 20 million transactions, so you could potentially be able to do this, but I don't know if you have sufficient power there. Overall I thought the paper was well executed with really nice empirical evidence and very impressive.

So well done.

(Applause.)

DR. LAIBSON: Given that we started late, we're right on time, but we don't have a lot of time for questions.

Do the authors have urgent responses that they want to make to the group?

Let's just jump right to a handful of quick questions, and I'll accumulate the questions, and then we'll just answer them in one fell sweep.

DR. FARRELL: I have a fundamental

observation which I think is prompted by Michael Grubb's paper, although I haven't read it yet. I plan to.

And that is that very generally, if you have market power or any situation where there's going to be a price above marginal cost, there is a total welfare gain in the model by bamboozling consumers into not paying attention to that price above marginal cost. And yet as an enforcement matter, I think it's crucial not to put too much weight on those welfare gains.

One question is: Is that's what's going on in your paper? And then a question for the group as a whole is: What's the right way for economists to say, "Yes, that's there, it's in models, but if you agree with me, we don't want to put very much weight on that from an enforcement or policy point of view?"

DR. LAIBSON: We're accumulating questions.

DR. INDERST: This is a very general comment to the issue of the overdraft payments. The presentations basically start with the idea that we basically know what's going on out there, that people are basically fooled, et cetera, and possibly something should be done. And there was a moment where I think one of these authors, he said if we save payments, how much it would bring us in terms of groceries we can buy.

But two very short comments on this, and they all kind of come back to the issue of pricing bundles. First, cross country evidence, if you look at different countries, even different European countries, you see a variety pricing plans for banks to make money with the retail customers.

So in some countries you pay a fixed amount for an account and in some countries nothing. Does that mean we have different regulations? Does this mean we have consumers with different neural stuff going on? So, how do you explain that?

Possibly there is a wide variety basically of how customers are charged, and it's not clear to me what could explain this.

Maybe inattentive consumers could be one reason, but to me it's not quite obvious why in some countries it could be more and in some countries less.

Related to this, a very short comment: If we don't think about the multiplicity out there, it's not just because we get a lot of solutions to documentation problems, but because there is an economic reason behind one of the other pricing models.

Why not start with this pricing? But for me, this pricing scheme where you charge less overdraft, is like an efficient enterprising scheme. You're pricing the most inelastic demand, so you price a bundle and you put the highest price on where the demand is elastic.

Of course the models didn't have to show much because there is the demand shown.

DR. LAIBSON: Anyone else? Great. No other hands. Let's go through the authors in the same order, beginning with Michael.

DR. GRUBB: To address Joe -- with the short time it didn't come out as well as it should in my talk. The reason for why penalty fees can be socially beneficial in the model of price discrimination is that they allow price discrimination with smaller distortion.

I'm not quite sure if it matches what you raise, but the idea is: If I can charge a low-based marginal charge and a high penalty fee in such a way that the expected marginal price is actually equal to marginal cost, it's actually efficient.

I've got more degrees of freedom. When people are attentive to get people to choose the efficient amount, the marginal price for ever unit has to be equal to marginal cost. When people are inattentive, I can have different prices for different units at different prices, so the expected marginal price is still marginal cost, so you still make efficient decisions.

But the extra freedom helps me sort people into different contracts, essentially because a high demand consumer, when you think of the expected marginal price, is higher because they're more likely to pay that penalty fee, but they're not choosing that contract, so that doesn't create an inefficiency. I'm not sure if that helps get at your question.

To Roman's point about the cross country variation, I think I don't have a good answer. That's something I would like to actually learn more about, the variation, and something I would love to try and understand better.

DR. ZINMAN: So in response to Roman's comment, one reason we punt on trying to say anything about welfare implications of the results of our consumer responses is we're aware that there's a supply response. And, indeed, the time path of bank pricing strategies in this market has been fascinating and is very much worthy of continued explanation.

The one thing I would say in terms of your notion of this: I was thinking about this in a Ramsey model and I think you would need some additional gyrations or richness to capture the fact that when people are exposed to something like a relatively subtle attention shock, that their demand changes. But other than that, I take the big picture point you're making as absolutely on point and something that merits further work.

MR. LAIBSON: Nicola, do you have anything to add?

DR. LACETERA: No, I think the comments were great and thank you.

DR. LAIBSON: Let's thank the authors again. When are we going to reconvene?

DR. ROTHSTEIN: 3:40.

(Whereupon, a brief recess was taken.)

PAPER SESSION THREE: CONSUMER CHOICE IN NEW MARKETS

FIONA SCOTT MORTON, CHAIRPERSON, YALE UNIVERSITY SCHOOL OF MANAGEMENT

PRESENTER: DIRK BERGEMANN, Yale University

DISCUSSANT: DAVID BALAN, FTC

PRESENTER: STEVEN PULLER, Texas A&M University

DISCUSSANT: TIM BRENNAN, University of Maryland, Baltimore County and Resources for the Future

PRESENTER: EUGENIO MIRAVETE, University of Texas at Austin

DISCUSSANT: JACK HOADLEY, The Health Policy Institute, Georgetown University

DR. SCOTT MORTON: We now have the session on Consumer Choice in New Markets, and we're starting with Dirk Bergmann from Yale University.

DR. BERGEMANN: Thank you and I'm very happy to be presenting. This is joint work with Alessandro Bonatti, and this work really begins with the observation that advertising on the Internet brings an important way for sharp views of advertising in more traditional media.

Traditional media is basically every viewer was listening to a particular program, maybe think about radio. Every television watcher was watching a particular program, and was faced with the same advertising, and conversely, advertisers have very limited ability to reach a particular audience with particular characteristics they might be interested in.

With the Internet, that's quite different. Just think of a search. When I put in a search for, say, mountain bikes, and clearly I'm revealing a lot of information that I'm interested in, what I might be interested in consuming, and many advertising firms basically share, to a large extent, some parts of the search, so think about this for advertising.

If I'm looking or if I'm reading a blog on bicycling, clearly there might be some inference to be made about my interest in bikes. And if my browser tracks my browser history, then again I can basically infer when I'm reaching the next advertisement or the next website, what my possible interests are.

So there's a sharp break in new abilities in terms of providing a match between a consumer, his preferences and then the advertiser. So, one view of the Internet is that basically society is going on the Internet and as it improves over time, as it increases the quality of the match between the consumer and the advertiser, we basically get a much better and better match between the interests of the consumers and the advertisers.

Of course although it has an impact which has been well documented in all the competing advertising media in the sense that there would be a shift away from the traditional media, which has less ability to target, to the Internet or other electronic media, which has a higher ability target.

I'm just giving you sort of a time series here, and what you see is that over the last few years, there has been a change, a dramatic change in terms of the composition of how advertisers place the advertisements across a different media, and therefore what we would like to do in this paper is try to first understand what the implications of targeting are in terms of how the advertisers allocate advertising across different media, in particular with regard to the implications for large advertisers versus small advertisers, and then second we want to understand what the role of targeting technologies is in the competition across different media.

So that's basically the outline, and in order to do that, we basically want to think about advertising as a matching process, as a matching process under substantial friction. Some of the friction comes clearly from the fact that as an advertiser, I may typically send out messages or I may place advertising where I may not reach the audience

I'm interested in. These are lost information, and then also there's a cost of duplicating. I might have the right audience, but trying to send the same message over and over again, that's something which I also want to avoid.

The idea of targets is basically that they allow me to reduce a friction, and therefore they lower the cost of the advertising. So the way I want to think about this, and to really think about the heterogeneity markets both in terms of customers being heterogeneous and also firms being heterogeneous, is that we want to think in a two dimensional world.

On one side are the advertising markets, I'm going to label them as A, and think about these as basically different media, different websites, different communication channels. On the other side, we have heterogeneous consumers and each one of them is basically just interested in a particular product, so they're single minded in the sense that they just want to buy a bicycle or they just want to buy roller blades or something like that,

We want now to think about the consumers basically being distributed across media markets, and also across product markets in a sense that that is dangerous, so a consumer is basically characterized by two dimensional times. One, what's the media where we can find him, where we can locate him, where we can reach him, and the other is: What is his interest, what is his preference, what can we convince him to buy?

You want to think about market structures in this world, and basically you want to think about the distribution. Although the consumers are on the one side, the media market, and on the other side is the preference, mainly what you would like to bring? This two dimensional representation of course then suggests that we may want to look at the cross-section.

That is, we want to say in a particular advertising market, and that is sort of the vertical lines, we are interested in what are consumers who are reading this media or following this newspaper interested in? What are their preferences? And likewise, as a seller of product X, that is as the firm that sells product X, I would like to know, and that's the horizontal line, where can I actually find my consumers, on which media markets do I need to be present and extend messages?

The basic exercise that we are going to pursue here is maintaining as given the consumer's preferences, which is this distribution of the products that they're interested in, which is S(X). Then we want to think about what happens if the slotting, that is the targeting of the consumers, is improved through the targeting technology? What impact will that have on the way you're going to place advertising, on the way we're going to price advertising, and on the social value of advertising in terms of matching consumers with firms?

So the two extremes are clearly in this scheme, the possibility of perfect targeting, a consumer with interest X is only found on one particular advertising market. That's the best situation from the point of view of the advertiser because he knows exactly where he can find his customers, and on the other hand, there is sort of the situation with zero targeting, when all the consumers, irrespective of their preferences, are in the same market or basically all the advertising markets show the same composition of consumers.

The question we're then going to pursue is: What is happening if you basically are moving from a world where we have very little targeting opportunities to one where we actually have a lot of targeting opportunities? That's of course just a semantic representation, but what you already see is that if you're going to have more targeting opportunities, then we will see that customers will move away from a few large media to basically smaller media or websites or blogs or whatever you want to think about. And at the same time, in each particular market, you will find a higher concentration of consumers with a particular interest.

So, in order to make this final, we are going to work in a frame work with dramatic distributions, that is the preferences of the consumers follow the dramatic distribution, so you think of the lambda as basically the point of concentration. Products X, which are nearby zero, are products which have a large audience, basically are mass market products which have a large X or small K. Products that have a large X are basically the long tail of the market which has a very small audience.

We also are going to assume a geometric distribution of the consumers across the advertising markets. So here on this scheme, what you're going to see in advertising Market 1, is we will find consumers which are interested in Product 1, but you might also find consumers which are interested in 2 and 3 and so on.

If you want to think in the language of

bicycles, think about sort of leisure bikes and then race bikes. Readers who are reading Sports Illustrated might be interested in leisure bikes or race bikes, but if we just say regular use, which is more focused towards sport bicycles, we'll only find guys who are interested in more specialized bicycles. So that's the sense in which these media markets are fine tuned in terms of reaching a smaller and smaller audience.

So, that's the model that we're going to work with. What we're going to think about is what happens when the targeting technology becomes better, that is, when the distribution to the consumers over this advertising market basically reaches or gets closer to a situation where we have perfect targeting, where there's a perfect match of the interests of the customers with the particular medium. So what I showed you is that when you think about targeting, then that has both a size and a competition effect on the market.

I mentioned that we want to think about advertising basically as matching, so think about this as random matching. If I'm going to send out messages as an advertiser to my population on a particular advertising market, I'm going to have a uniform probability of reaching a particular customer. And so if I'm sending out more messages, I have a higher probability of advising my customer in a sense that I'm going to make a sale, but of course I face the risk of redundancy if I reach the same customer twice, and that's something I want to avoid.

My ultimate advertising policy in each one of these advertising markets is basically a policy which tries to resolve it.

Now we're ready to describe the equilibrium of these markets and advertising policy for a firm is not to actually decide in which market it wants to be present in and how many of the messages. It's what volume of advertising he wants to achieve in these markets.

And so we can say, "Well, the gross revenue is simply coming from the number of matches that I can generate and support by my advertising policy," and what I want to resolve in my policies is basically the optimal trade off.

With the geometric distribution, we can actually solve for the demands in the competitive equilibrium and we get a linear demand function, so we can then trace out the impacts of a better targeting technology on the markets.

Another nice feature of the exponential distribution is that we have a certain station area that the markets basically -- whether or not it is the same in terms of the competition, they all look the same with respect to the relative composition in each of these markets, so the prices are actually equal across all of these markets.

So what do we find? As soon as there was some targeting opportunity that is much larger than zero, we'll find that both small and large ones will advertise. And that's a fact that we'll be presenting on either the small or large market, depending on the size. And so now we can ask: First of all, what is the social value of targeting? What are the numbers of matches that are supporting advertising technologies, which have better targeting opportunities?

And here we find that it's uniformly increasing the wealth here simply because it supports more matches. The question then that we ask is: Well, how was that increase in the social value supported in terms of the purchases that the consumers are actually making?

What we see is that we see a change in the advertising policy in the sense that firms start to purchase less volume, and it's the smaller firms which take up some of the slack, so this is the long tail and Anderson. But it's also going to be true that the number of participating firms in each one of the advertising markets is declining because in some sense the interest in any particular advertising market focuses on a smaller and smaller segment of firms or products.

So that leads us to ask: As we increase the value of the social technology, will the advertisers be able, over the media, to cover or recover some of the gains in terms of the social value of the matching in terms of the price?

What we find here is that initially as we increase the targeting possibilities, indeed the price for the advertising is going up simply because we have a higher value that we're offering per a message being sent, but eventually then the price for each individual message that I'm going to buy is going to go down in the equilibrium. And the idea is here that as the messages become more and more efficient, the concern of the advertisers is moving away from not hitting the right people to sending the same message often to the same people.

As you move from the cost of not reaching the right guys to basically the cost of saturation or duplication, we will also see that the price is changing and then returning as it's decreasing, even though the value of the advertising technology is increasing.

There's some interesting empirical evidence in the sense that in a recent paper, Chandra and Kaiser looked at the value of targeting in the magazine market, and there they find that magazines that can offer more homogenous readership are able to extract a higher value.

On the other hand, if you go to the Internet markets, where there is a point to be made that targeting effect is already very high, Rutz and Bucklin show that branded key words -- that is, key words which focus even more narrowly on a particular search -- obtain a lower price on the search engines. And the same is also true for longer key words because they are also targeting finer and finer markets, for which in some sense there's less interest and less competition in reaching the customers.

This is basically the first set of results that we pursued. What is the implications of targeting for advertising? And the second step to take is basically to ask: What happens if we move away from the single homing that was implicit in the first part of the paper to a world of dual homing where consumers can be present both in Media 1 and Media 2? And of course in our world, we're interested in whether it can be present there online as well as offline.

So the relevant comparative study is to explain what happens if the consumers spend more time, hence can reach more messages on the Internet, that is they've increased the presence of the consumers on the Internet relative to the time they spent on the traditional media.

So, we now have competing media to do the dual homing, and we can ask: What were the prices in the offline media and what are the prices in the online media? What's the distribution of the consumers across these two media, and what happens when consumers spend more time on the Internet?

Clearly what we find is that will always be the case, this is that only the large firms can reach a large audience which will be present both offline and online in terms of the advertising policy. The small firm will not go into the large media. They will always be trying to get into the targeted media.

What's more interesting is to ask: What happens to the price of the offline media or the revenue of the offline media when we are be moving away from the offline media to the online media?

Here what we find is that the price that the offline media can distract from the advertisers is decreasing linearly in the presence of the viewers online, whereas if I were to think of a second competitor offline, the decrease in the price would only be to the square of the presence of the competitor.

What this then shows is that the merchants of the Internet have a disproportionately large effect on the revenue of the offline media in terms of their advertising revenue, and that's driven by the ability of the Internet to provide a higher value service in terms of its targeting ability.

That explains some of the decline in the revenue of the traditional media, and in particular those traditional media which can't, by construction in some instances, target as finely as the Internet, and here in particular the newspaper.

Again there's some interesting evidence to that contest, that there is indeed competition between online and offline. But perhaps I should just mention that here we tried to present the model that allows us to think cohesively about heterogeneous consumers, heterogeneous advertisers, and therefore the possibility of improving targeting technologies. But of course our view, say of the media, in particular, is rather limited in the sense that we were only interested in the revenue coming from the advertisers.

We might also be interested in thinking of a dual market where in some sense the media has tried to get revenue both from the readers as well as from the advertisers. And these are logical next steps to take. Thank you.

DR. SCOTT MORTON: Thank you.

DR. BALAN: So I am Dave Balan. I work here at the FTC. These views are mine, not anybody else's.

This paper is about informative advertising, and you have some consumers and some products, and it's a good thing when the consumers learn about the products that they might want to buy, so welfare becomes pretty straightforward. It's a good thing when the right people find the right price.

So, the authors have a whole bunch of environments that they model. This is a very elaborate thing. I have market in square quotes, because this is the FTC, and we mean something different by markets, but here we have a single advertising market and one medium, so that's like a world where there are newspapers, and there's only one newspaper.

The second one is a world where you have newspapers, and then you have two media likes newspaper and TV, but they're the same size. Then you have newspaper and TV and they're different sizes and then you have newspaper and TV, and they're different types in a way that I don't have time to go into. And then finally you get to online: one online, one offline, and then many online media, which I think is the one that's ultimately of the most interest.

So there are a whole bunch of parameters floating around in the model. The three biggest ones I think are we've got this lambda parameter, which is the concentration in the market, again in square quotes. When lambda is big, lots of people want a small number of products. When lambda is smaller, they're a lot more diffuse in what products they like.

Gamma is this targeting parameter that Dirk was talking about, how easy it is to identify in the one pole where everyone is in the same place and in the other pole where exactly these kind of people are in exactly this market so you know exactly where you have to go to reach them.

And social welfare is unambiguously increasing enough because this is all about matches and it being a good thing when matches are made.

When you have the online and offline, beta is the fraction of the time you spend on the Internet instead of watching TV or reading the newspapers or whatever.

Then there's a whole bunch -- and I mean a whole bunch -- of comparative static exercises on parameters and some others, the effects of these parameters on advertising prices, advertising prices per person successfully reached, who in equilibrium which firm sizes advertise on which media and which consumers pay attention to which media, so just a whole bunch of stuff.

Then the comparative statics are not all simple. They're like all these non-monotoicities, so I'm going to try one, and I think I got it, but this is lambda. That's that concentration parameter and in the first environment, that's one newspaper and there's only one newspaper, so the simplest environment. When lambda is low, which means that people are diffuse in which products they like, you increase lambda a little bit and what happens?

What happens is everybody who was already advertising their market share -- although I don't know if that word just means the base of people who are potentially interested in their product -- goes up, so their demand for advertising goes up. So the price goes up, but if lambda is already high and you make it a little bit bigger, then lambda goes up. The marginal firm actually has a smaller base of people who are interested in this product, which seems to me the demand will go down.

But on the other hand, a bunch of the infra-marginal guys have the number of people interested in their product go up, which tends to push it up. But on the other hand, there's diminishing returns because if you hit somebody twice, that doesn't help, and the net effect of that is, when all the dust settles, that it's negative.

This is an impressive undertaking. There's just a great deal here, and the results are built up step by step in this manner that was described in a very appealing way. The downside -- I don't know how to finesse this -- is the online versus offline results are on page 28 of the 34 page paper, so I don't know exactly how to handle this. And there's so many results that it seems like some sort of targeting of which ones y would be of value, but that's just because there's so much there.

It's bad form for a discussant to talk about the paper that they wish you wrote, but I'm going to do just a little bit of that. This is about informative advertising, and it's not about informative advertising of prices which we know what there's lots of. You open the newspaper and then you find out that products you already know about are on sale.

This is learning about the existence of a product, which surely happens. Certainly new products are being introduced. That's how they get people to know about them, and in the example, it wasn't in the paper but it was in the presentation. Maybe a good example of learning about products are these things like when you like a book and then Amazon tells you a book that you might also like, which definitely has the flavor of informative advertising because it's really something you wouldn't have otherwise known about.

But I hadn't thought about that before, which weakens the force of this comment, but I'm really somewhat skeptical of how often this is really about product existence, and in the paper -- and this is a model, an assumption I am showing you -- I wouldn't want you to put too much weight on this. But in the paper, if you don't learn about your best product, you buy nothing. You don't buy the next best product that you know about, so I'm skeptical of how often it is that we have a world where people are at sea until they see an advertisement and then it's like, "Ah-ha, now I know about this thing that will bring me inner happiness."

And moreover -- and this I was worried about, as I understand it, and I could be wrong -- the model says that people look at media to learn about products. It's not that they look at media for content and then the advertisements are piggybacked on to the contents. And if I'm right about that, then the bicycles example -- which I'm glad Dirk mentioned in his presentation -- he has this idea that there are three kinds of bicycles. It's like cheap old toy bicycles, decent bicycles, and serious bicyclist bicycles.

And the serious bicyclists might read Bicycle Enthusiast Monthly and Sports Illustrated and The New York Times, but a guy who just wants an okay bicycle reads Sports Illustrated and The New York Times, and the guys that wants a cheap bicycle will read The New York Times.

That might be true if they were reading these media for content. If they're reading these media to learn about products, then I don't think the bicycle enthusiast is going to read The New York Times because he knows he's not going to see an advertisement for the fancy bike that he's interested in The New York Times. So I wasn't exactly sure what to make of that or exactly how many bodies are buried there but it was a concern.

But the biggest thing -- and this is me, and you won't be the first ones to think I'm crazy, if you think I am crazy -- is I think advertising is about persuasion. Persuasive advertising is what it is about.

There is a paper that says one quarter of GDP is persuasion. They did this really super cheesy calculation which is not to be taken that seriously but a lot, a lot, a lot of GDP is persuasion. But if I was the social planner directing talent into problems -- and I think pretty massive amounts of talent that were brought to bear on this problem -- I would redirect to things related to that problem. But that said, that doesn't mean this problem is not worth thinking about.

So a very, very rich, impressive thing, carefully and logically derived; the very scale of it makes it a little bit difficult to digest, and there might be some tweaks that could be made in directing a reader with finite time to what's most important. And then this comment, which you can make of what you will.

Obviously you have this paper now, but I think figuring out how persuasion works and what the welfare implications of persuasion are is kind of where it is at in advertising.

That's it for me.

DR. SCOTT MORTON: Great. Well, that sounds good. We'll move to the next paper. Let's thank the author and the discussant.

(Applause.)

DR. PULLER: In this session on new markets, the new market I'm going to talk about is essentially what happens when you allow homeowners to choose who their electricity provider is. We're going to analyze what happens when you allow retail choice in residential electricity markets, and this is a joint paper with Ali Horatacsu and Seyed Ali Madanizadeh.

Texas, like a variety of states, has moved from one means to another means of procurement of residential electricity, so from a regime where everybody will buy from a regulated incumbent for a regulated rate to one where you can choose your providers, and those providers might have discretion as to what prices they charge.

And a lot of the motivation for this is retail competition might add value-added services or if a lot of consumers are actively searching, that might make the market more competitive. So the models in these markets often give people the power to choose.

Just to give you a little hint as to what our findings are going to be, we have an alternative title, which we haven't quite had the guts to put in the published version of the paper, but this is probably a more apt title for our paper.

I'll give you details of how retail choice works in a second, but briefly starting in January 1 of 2002, all residential customers in Texas were assigned by default to a retailer that was basically affiliated with the incumbent, and then every month consumers could potentially switch around from that incumbent or back to that incumbent for a variety of competitive retailers.

What this graph does is it shows the evolution of market shares over the first about four years of the market. As you can see, there's a gradual erosion of the market share of the incumbent. You can see there are two providers that have about 15 to 20 percent of market share, and there are a couple of smaller providers, and then a variety of actually even smaller providers that I didn't even put on the graph.

So, here's a question for you: If you had to guess just based on this figure what rates were charged by the different retailers, what would you guess? You would probably guess the incumbent has a higher rate, and indeed that is true, but the question is: How much higher? And that probably depends on your perceptions of the savviness of Texas home owners.

So, here's the answer. This is an average rate for a typical usage amount. The line in blue here with the dots is the average rate for the incumbent, and all the other lines are rates for a variety of other competitive retailers.

As you can see, except for a couple of months, there's at least one and sometimes more than one competitive retailer that has a rate, one to one and half cents, sometimes even little more than that cheaper than the incumbent.

So this motivates us to ask why, and if you aggregate this across a month or across a year, there's been substantial savings there. And then why, despite potential savings, does there seem to be a fair amount of consumer inertia keeping customers with the incumbent?

What are possible causes of consumer inertia? We are going to try to group them into three categories and then try and empirically quantify these. So the first category could be that, in fact, electricity is not a homogeneous product. Anecdotal evidence is that people perceive that the reliability of their power is a function of who their retailer is.

Technically that's not true at all. It's still the same power, runs at the same volume, same meters and everything else, but maybe people have the perception that if there's a power outage, it matters who their retailer is, even though it doesn't.

The second category could be that maybe customers suffer from status quo bias, and they just don't pay attention. We're going to call this decision cost. And the third category is that maybe there are actual switching costs. For example, non-monetary switching cost might be that I have to get accustomed to what a new bill looks like and change my online bill pay, for example.

We think there are merits to trying to quantify these effects, in general because these might show up in other retail choice markets, but for this particular case, because there could be policy implications. So for example, if I think that it's kind of a brand effect, maybe that will erode over time, and if so, maybe you just think about that as kind of a transition cost to retail competition.

If they're decision costs, maybe public information campaigns can reduce those costs. Are they switching costs? I'm not really sure if there are policy levers to influence that. Maybe there are, maybe there are not.

So the goals of this project are to quantify those and to see if there's evidence of product differentiation, search cost and switching cost, and then also to ask the question: Are there heterogeneities across different demographic groups? So loosely, do different demographic groups benefit more or less from having retail choice in residential electricity?

This is related to a variety of literature, works that have looked at retail choice and things that are traditionally utilities, whether natural gas, telecomm or electricity, but more generally to a literature that looks at what happens if you take some product and move from a regime where everybody has to buy that product at some regulated provider -- there's no choice -- to a regime where there is choice. And that product could be a school. It could be health insurance. It could be long distance telecomm and asking what the distributional consequences are of adding that choice.

So what I will do is just give you some basic descriptive statistics that look at the raw data, and we'll build a model, which is going to econometrically be able to test for these three different effects, sources of inertia. Basically we're going to find that there is an incumbent brand advantage, but it tends to erode over time.

We're going to see that people don't search very much, but there is a seasonal pattern, which I bet you guys can guess if you've ever been to Texas a certain time of the year. And we're also going to find that there's some demographic heterogeneity, and I'll tell you what that's going to be.

Some of us might have personal experience with retail choice. In any state that's not dark blue, there's been some experimentation with choice either at the residential level or the commercial and industrial level.

The way it works in Texas starting in January of 2002, is customers were assigned to a retailer that was affiliated with the old incumbent, so an AREP. That AREP had a required rate called the price to beat, which was actually a cut from what it had been before, but because of a variety of things that were going on in the wholesale market, that was still thought to be above competitive levels.

Regulators thought through some logic that that was a good thing because it basically guaranteed head room for retailers to enter into the market. The rate could be adjusted over time, but it was adjusted in a way that was indexed to cost.

So for the competitive retailers, these are just kind of intermediaries that procure power from generators and then market those to customers. It turned out the largest CREPs were also AREPs from other parts of the state, so brand names that were known in other parts of the state. By the end of our sample period, which was 2006, there were typically over ten CREPs.

So critical for understanding why there might be consumer inertia is finding out how do consumers get their information, so there are media sources, as you guys can probably guess. One critical website that was actually created by the Public Utility Commission was basically viewed as a one stop shop to learn about rates and to conduct a switch.

Here's a screen shot of what that looks like. It may be a little small, but basically what you do is you type in your Zip Code, and then it will give you a list of a variety of different providers and characteristics which you can sort. For example, they will tell you what the average rate is of a typical usage level of a kilowatt hour, and it will tell you what a thousand will cost, just in case you can't multiply by a thousand. It will tell you what the rate structure is, whether it's variable or fixed, whether there's some wind energy blended into that, and what the terms and cancelation fee are. Then if you want to switch, you can click through and find out other information and actually click through and conduct your switch.

What we're going to look at is a certain service territory which you can see in red here. It was spread across the state because it was formed as a result of mergers over time. It's nice for us because we're going to get some nice demographic mix of urban, suburban, and west Texas rural areas. We're looking at the first four years of the market and about 200,000 customers.

If each of those customers were actually technically metered, what we have is each month who their provider was and how much they consumed that month. Then we have the address of the meter so that we can link that to characteristics of the census block groups, not on an individual level, but census block group characteristics.

And for each retailer, we have the rate plans that were offered. So what we can do with maybe a little bit of error is calculate what the bills were and counterfactually what the bill would have been if they had bought the same power from any of the other providers. And we're going to focus on the six largest retailers.

Just to give you a little descriptive data before I go into the model, this just shows a count of the number of switches for each month in our sample period, so as you can see in the first year, there wasn't as much switching behavior as there was in later years, and if you squint but just a little bit, you can see that there's a peak every year, and that peak is going to correspond to June, July and August.

So now we described what potential savings could be, and I just mean this purely as descriptive evidence. We're going to ask the question: For households in the months that they're buying from the AREP, the incumbent, what would their bill have been if they had bought the same amount of power from any of the CREPs?

Now, this isn't a welfare calculation by any means. It's not accounting for switching costs. It's assuming that they forecast everything right, but it will give us some idea of the magnitude of what potential savings would be.

To do that, we're going to have to make some assumption about how savvy consumers are and how often they search, and we're going to have two extremes. At one extreme we're going to let these guys counterfactually switch only once, and they're going to switch at the very beginning of retail choice to one of the other big named providers, one of the big named CREPs, and depending on who they switch to, it's going to be somewhere around $7.50 to about $10.

At the other extreme, we're going to imagine that every single month they switch to the lowest cost CREP, which is technically infeasible but it will provide us an upper bound, and then we're getting members who are at a little over $12.

So we think somewhere between 7.50 and $12, and just for basis of comparison, other energy policy according to CBO would have cost each household a little more than that but not that much more.

Finally, descriptive evidence about the demographic effects. We would like to say which demographic groups, or at least neighbors, might take more advantage of choice, so to do that what we're going to do is calculate household level metric of percent achieved. We're going to calculate when they're with the incumbent, what the bill is with the incumbent, what the bill counterfactually would have been if they had bought the same amount of power from the lowest priced CREP and then what their actual bill is, and then we're going to calculate this measure of what percent of potential savings was actually achieved.

So this is going to be a household-month level observation where this variable is going to be something between zero and one, and then we're just going to regress this on census block group demographic characteristics, again just descriptive analysis. Again this is an individual level data, right, and there's heterogeneity within a block group.

I encourage you to think about this more as neighborhoods as opposed to individual level effects, and we're seeing that there's a high percent of realized savings in neighborhoods that have more college educated, more African Americans, fewer Hispanic, fewer seniors and a lower poverty rate.

Now, I'm not going to view that as a demand side estimate because obviously some of these demographics could be associated with things that firms might do, like advertising. So I'm not necessarily viewing this as the demand side, but it's still kind of characterizing which neighborhoods seem to be taking more advantage of choice.

Now let me go into the model. We're going to imagine choice as every single month for households engaged in a two-stage process. The first stage is they're going to basically decide whether or not to look around.

If they don't look around, they're going to stay with the same provider, which we're going to call Provider K. If they do look around, they're going to engage in a standard discrete choice decision. They're going to observe the product characteristics of the alternative retailers and choose the one that maximizes utility, keeping in mind that they can decide to stay with their existing provider, Provider K. And when we estimate this, we're going to allow for heterogeneity across both the probabilities of looking around and the choice probabilities.

One set that's going to help us with identification is institutionally there's a group when you move into a new house, you don't get any power until you make a decision. You have to actively choose, and so we're going to model these people's deciding with probability one.

In the paper there's a formal model, but let me just give you a simplified model to give you a sense of how we're doing this. So, let's assume there are only three retailers. Everybody is the same, and us as analysts only observed two months of data, last month and next month.

Let's assume that each household currently with Retailer K is going to search with some probability that is specific to that retailer, and then those that decide to look around, conditioned on deciding which households are going to choose the Retailer J with probability P(J).

In the simple model, we have five parameters to estimate. The probability of looking around for each of the three retailers, and the probabilities of choice: P(1), P(2), and then P(3) is just going to be P(1) minus P(2). We've got five parameters we want to estimate, so how can we get that?

What we're going to do is create a matrix that's counting the number of people as a function of who their last provider was in the rows and who their next provider is in the columns, so since people don't switch that much, the diagonals are going to be well populated, and the off diagonals are going to be less populated.

We can write down an expectation of the numbers that are going to be in each of these cells. Let's assume that the total number of people that have been with Provider 1 is N(1). In this first cell, in expectation, we would see that for each individual, the probability that someone doesn't decide to look around plus the probability that someone does decide to look around but then chooses their current utility.

In the second cell there is simply the probability that someone that looks around times the probability, conditional on looking around, and then they choose Provider 2, and the same thing for number three.

What this matrix is going to give us is nine moments. Now, it turns out that in each row, one of the moments is redundant, so we're only really getting six, but we can use these six moments to estimate the five probabilities. So in simplified terms, that's basically what we're doing.

Now, we want to interpret these lambdas and Ps, so to do that, we're going to actually parameterize both of those things, the looking around function and the choice function.

As far as the choice function, the lambda, we're going to make that just a nice kind of S shaped function with the variables that are going to enter into a retail variable, so who the last retailer was, the month of the year to allow for seasonality and searching, and in some of our specifications, census block group demographics.

And for the probability function where you use a standard discrete choice set up, we're going to use a logit model, where our product characteristics are going to be price, a dummy for the incumbent to pick up an incumbent brand advantage, that interactive with months to allow for that brand effect to vary with time, and then to identify switching costs, we're basically going to have a product characteristic that is, "I don't incur switching costs."

So who doesn't incur switching costs? You don't incur switching costs if you're staying with the same provider and you're not a mover. Now, it turns out that the identification here is really coming off a nonlinearity in the logit probability so it's a functional form assumption. So if for some reason you're not comfortable with that, we'll present results not estimating switching costs and then estimating switching costs.

We're just going to pack all this into a GMM estimator. Think back to that three by three matrix; a count in each of the cells is just going to be equal to the probability of anybody being in that cell.

I'm going to show you results that are in a couple tables like this. Let me be clear on what's in these tables. The top two panels are parameters for estimating the choice step and the decision step, so it's just parameter estimates. And then based on those parameters, we can calculate things that we can interpret. So mainly, we can calculate the probability that a customer looks around as a function of which retailer they are with.

And then given that they're looking around, they're deciding the choice probability, and then we can turn and also calculate what the price elasticities are.

What we're finding here in our basic results where we're not estimating the switching costs is that, in fact, there is a brand advantage, but it tends to erode over time. In terms of the estimated price elasticity, we're finding that the regulated incumbent has fairly inelastic demand and then the smaller competitive retailers have a demand elasticity estimate more around five.

In terms of when people search -- these are the dummy variables of seasonality that are entering into that lambda function -- we're finding that people tend to search more in the summer, particularly in July, but that being said, they don't search that much. Depending on the retailer, they're switching anywhere from 2 percent and 5 percent of the months. No estimates switching costs, which again is identified off this nonlinearity in the logit probability, we're finding evidence that there are non-trivial switching costs.

We are still doing a robustness test to test for the appropriateness of the model, but let me show you some preliminary results where we are allowing these effects to vary by demographics. We're getting results that I think are at least roughly consistent with those reduced form regressions I was showing you before in demographic characteristics.

So, what we've done here is we've taken the price coefficient and the incumbent brand advantage coefficient and interacted that with the demographics. Again we're thinking neighborhoods here and we're finding that you're getting more price sensitivity in neighborhoods that have more African Americans and more college educated people.

You're getting the brand advantage being lower in neighborhoods that have more seniors, more African Americans and more college educated people. And with the exception of the seniors, this is qualitatively consistent with the reduced form results that I was showing you before.

So conclusions: We're finding that you if you entirely homogenize the product, we're getting dollar figures of $7 to $12 left on the table per month. Our model is suggesting evidence of all three of those sources of consumer inertia.

Our next step it to do welfare calculations and we're interested in thinking about this incumbent-brand effect. Let's imagine that the major driver of that is this perception that power is more reliable, even though we all know that's not true. Should that count in welfare or not? And maybe this is a philosophical question.

We're going to show results of it both ways, but the numbers are going to be very different I conjecture, and I guess readers will be left to interpret whether that should go into welfare calculations or not.

Thanks.

DR. SCOTT MORTON: Thank you. Tim Brennan is our discussant.

DR. BRENNAN: Thanks for inviting me here. We learned this morning that disclosure is important, so let me disclose four things. First, I'm an antitrust person, not a consumer protection person. Second, I'm a neoclassical Neanderthal, not a behavioral post mori most.

The third is that I probably know less econometrics than anybody else in the room, and the fourth is that I'm over the Laibson hill, so one may wonder why I'm discussing this, and I may be one of those soon.

So overall points I want to make here: First, I don't think that these results are particularly surprising. I'll talk about that a little bit. I think that the main thing that I took out of this is the accomplishment, I guess, and the thing that makes me think the most about this is how it unpacks brand loyalty, switching costs, and search.

I want to talk a little bit about how this gets framed in some of the papers today I think about. Is this behavioral or is this response to cost? I'll mention some policy implications -- and Steve talked about this a little bit about what difference it makes -- and then I want to talk about something which he didn't talk about in his presentation, which is that his measure of pricing was average, not marginal prices, which actually he makes a big deal about in the paper, and I don't disagree with that, but I wonder what that means.

I'm going to just blast through this. As he pointed out, there's been a lot of work on residential electricity choices. There's lots of information on reluctance to choose. I should point out that Texas, which after four years had two-thirds to three-fifths of the people still using the incumbent, is viewed as the great success of residential choice.

In most places it's in the single digits in the States for one reason or another, so Texas is actually quite unusual in that regard. There's been extensive efforts at persuasion. He talked about some of those, and as we've heard a little bit earlier today, the choosing not to choose is not unique to electricity. I mean, how often do we change our brand of toothpaste or cereal or this, that or the other? And the idea that people stick with what they've got is really not that unique to this.

Just to illustrate, this is something I think I've shown here before; this is from a Pennsylvania guide to help people choose their electricity provider. It's a little hard to read, but it has eight little tips here and things about each electricity generation supplier for you to save, the price would be this, multiply something on line 3 by line 4, divide the subtotal.

Then we have a picture of a family, and we can't really see it very well, but it looks like they're really happy on Christmas day, opening an envelope saying that Santa has given them the right to choose their electricity supplier, and they're delirious.

These are the questions to ask your electricity generation supplier. This is a woman who is happy, I guess just before she puts the gun in her mouth, about having to ask some retailer all of these things. So I think, "Why are you bothering to do this?" It shouldn't be a surprise that anyone wants to bother with this. Wasn't it fine before? And a lot of people who aren't economists ask that about opening these markets.

Now, part of the reason that I found this paper so provocative is that if you think about choosing to stick with the incumbent, there's brand preference or switching costs or search costs. And it's hard even for me to tell those about just thinking about them, much less econometrically as differentiation risk aversion or just not wanting to bother.

Someone mentioned Schmalensee's pioneering brands earlier but I think that the risk isn't reliability risk. I think it's a business and price risk. Here is this company that's been around for decades and here's some fly by night electricity provider. How do I know they're going to still be in business in six months?

I think that's the risk, not whether there's going to be an outage or not particularly. If search is seeing that the new provider is as good as the old in terms of being reliable in that sense, can one separate brand preference from search costs? After all, what you're searching is to find out if they're like the incumbent or not.

If switching costs are low, does that reduce the cost of search, at least when we're talking about ex ante verification about that? As I was thinking about this, I was wondering whether Steve could appear in those old Miller Lite ads where people are talking about the demand for Miller Lite beer where it was like "More taste," "No, it's less filling," "No, it's more taste," "No, it's less filling."

They have these arguments about it, and actually it may be a career for him because I was trying to Google it and see if I could find a picture of an old Miller Lite ad on the Internet. I couldn't find one. But I did find that Miller is going to revive that ad campaign, so Steve can be sitting there in the bar, and when the argument breaks out, he can go up to Peyton Manning and say, "Look, I have this generalized method of moments regression; it will tell you whether it's more taste or less filing and how much." So that's kind of how I saw this.

did he and his coauthors do it? I'm not sure about this because I'm not sure that what he did matched the simple example for reasons that I'm going to indicate here.

I'm going to talk about the 3 by 3 matrix that he had up. It looks like you have N brand preference variables, and then you had the probability of switching. And you had that P(1) P(2) there -- and at least in that diagram -- that probability of switching is independent from where you started from. That's what gave you more moments than values and so that's how that would work.

I wonder whether that's true in the sense that if you went to somebody because they were green, I might be more likely to pick an alternate green provider. Someone who is making those switches just may be different, and if it is different, then I'm not sure that you've got the extra degrees of freedom or what exactly you want to call it, to be able to pull this out.

One of the things that I liked about this is that it is an attempt to characterize this in cost or preference terms. And I would take how much people are losing a month by this, which is an impressively large number, as basically their willingness to pay to avoid being put in the position of that poor family in Pennsylvania.

What difference does it make? It depends upon the policy objective. If you're going to make these markets work come hell or high water, then it depends. If you think it's search, fix this. If you think it's switching, fix that, and so on, like unreportability in telephones for example. But if the question is whether to have markets at all, it's all part of the same cost basically.

Maybe the incumbent would be free to exploit that, and someone that switched may be worse off because maybe they preferred being in the regulated environment. I might just say rationally, let the Public Service Commission do it for me, I've got better things to do, and let the 65 percent of the market, commercial and industrial, make the choices as actually they do.

Finally, just a couple things on something Steve didn't talk about, so I'll be very, very brief about this. He actually says -- and I don't doubt this -- that the relevant price he wants to use is the average price people pay, not the marginal price people pay for electricity. There's a lot of data on this, and actually the UCEI released a paper I think yesterday. This is a common argument in electricity.

And it may very well be. I'm not going to give some sort of neoclassical, "Wow, this is real. What's going on about it, but what does it mean?" If it means something like on the third bullet there, looking at the fact that if price is constant, it doesn't matter. But if there is some variation, if it really is maximizing the value minus the average revenue times output, that's value minus revenue and that assumes that they're even more sophisticated, which probably is not it.

If that's not it, then you need two equations. One is that there's this price that they take to be constant. The second is that price happens to be equal to the average price, and the problem with that is, I think, you get multiple equilibria, so I don't know exactly how that plays out.

And part of the reason I want to know this -- and this is the last slide -- is to ask a question like what Joe asked before, which is: I'm intrigued with this average pricing idea because I don't know what to do with it if I was in a market where I really thought this is what people were doing. What does the SSNIP mean, putting my antitrust hat back on here? What does the upper pricing pressure mean? Where are the models that tell us what somebody is going to do if they're in an average price world? What would a monopoly do if they thought its customers were acting that way?

I'm not exactly sure what the answers to these are. Maybe if I had a little more time, I might come up with something, but maybe people already have answers for this, which is fine.

That leaves a question about how we do cost benefit tests generally if the area under a demand curve is based upon average rather than marginal prices. What does consumer surplus mean in that context? Those are some of the thoughts that occurred to me reading this really great paper.

So thanks very much.

(Discussion off the record.)

DR. SCOTT MORTON: Thank you.

DR. EUGENIO MIRAVETE: This is joint work with Jonathan Ketcham and Claudio Lucarelli and Chris Roebuck.

Thanks for having this paper. It gives us a great opportunity to revise it, and actually there's already a new version available on my website. We're going to talk about Part D.

Part D started in 2006. Oh, the disclaimer, last time I presented this paper, it took me over two hours, so 18 minutes is out of the question. Many of the details are in the paper, but I won't be able to go over it.

Participation rates are over 90 percent. Part D has expanded prescription drug use and lowered out of pocket drug prices, so beneficiaries are generally satisfied with that. The cost of the program, which exceeds about $39 billion, is less than what initially was thought.

So the question is: If everything is so great, what is the controversy about it? So we think it's whether consumers are making the right choices and if choosing among many different plans is beneficial or not.

Just to give an explanation for the title we have, this is the presidential address of McFadden seven days after the program was implemented, and his judgment was that the program was not going to be a success. He, in later papers, changed that view. A few months later Krugman thought that seniors were making choices among many options. They were of course seeing it as very complicated.

Then we have Thaler and Sunstein offering the explanation that people have 46 choices, and telling consumers anything is like not having help at all. They haven't updated this reference.

Another comment about health insurance, so all this gives you an idea of what the initial consensus in this market is. People were very skeptical about consumers being able to make choices that minimize the costs or direct consumption.

We have more evidence. All these are references referred to in Medicare Part D. Many of these papers are making use of the survey data. Some of them are not, but in essence, they go in the direction that consumers, one way or the other, are making mistakes. They have some bias because they place more value on certain features.

There's a paper by Abaluck and Gruber, valuing some features like monthly fees other than the prices of the drug and so on and so forth; they could perhaps question where many of these things are coming from a cross-section or lab. I have nothing against that, but we cannot look at what is the effect of the market anyway.

Here are some more of the papers related to other industries, so on and so forth. I don't think I have to go here, but we could perhaps start educating ourselves when we are going to make the decision on Medicare, you have to make a choice among several different plans.

There is a total of 39 markets, and essentially you have to sign up for a plan. That's in those payments every month, and you get some sort of discount on the drugs. The plans are very different. They cover different medical conditions. They have what sometimes are hated features like this donut hole, so coverage stops if you go beyond about $2,200, and then restarts later on at $5,000.

You can assure against them that this is initially decided -- a sign to give an incentive to consumers to reduce consumption or take care of consumption. Now, let's try to get a little bit more into what the problem is from the choice point of view.

So consumers, and we're talking about the elderly here with cognitive problems, have to choose many times between 50 different plans. The choice is made at the end of the year, then there are six weeks of an enrollment period, and you stick to that choice for the rest of the year, except if you are a low income individual. We are not looking at this in this paper, if you change to different a market.

Some information you can gather from websites, but the information about the choices in 2006 didn't exist that much. Then give an incentive for the people to sign up for the program if you think your premiums will go up 1 percent every month. This is an extremely interesting obligation. It's a unique opportunity. Everybody in 2006 is the same. They have the same experience independently of age. So in a sense we can separate what the effect of aging and experience is and avoid issues like state dependence and so on in terms of the estimation.

The next one is a big one. We go and look at 2006. We figure out that some people make mistakes, so on and so forth. People maybe have some bias for one or the other. What do we conclude about this?

Should we help consumers or are they rationale? What kind of value do they have? These decisions are repeated over time. So that opens the possibility of learning and perhaps switching them into different plans. Essentially this is the paper.

Here we have this vision of overspending. We have data for 2006 and 2007. Here we have the brief distribution of overpayment defined as whatever you consume. Now look at the alternative, 49 different plans, and let's see how much you are paying.

So here the mean is about $550 in 2006.

UNIDENTIFIED SPEAKER: This is overpaid consumption and you take the minimum?

DR. MIRAVETE: Yes, actual to minimum, including the possibility of no insurance.

In 2006, we have lots of people around $300 to %500, and then the distribution obviously switches in 2007. And lots of people are now in these smaller confines of about a $100.

This could be for a variety of reasons. You could say, "Wow, it's just learning. Everybody is switching, and everybody is becoming very smart." It could be also the fact that the plans have been changing over the years, and now there are less potential gains. But the shift is still there.

The distributions like this, the differences you can see between the mean and the media, are shown by a line there. Some people actually pay a lot of money, much more than other plans. So what are we going to do? Well, essentially we're looking at all the data we have, just trying to figure out what we can learn.

This is unconditional of everything. And the rest of the paper is pretty much looking at: Well, what can we learn? Can we learn something about it? Is there switching behind this? What exactly do we have?

We have a huge data set. We are working here with about 75,000 individuals. We have searched all their consumption. They are in these 39 different markets. Remember there are no low-income individuals because they have a different regime. So, we know some things. We know their age. We know gender. We have information about their health status. We know what drugs they are consuming, so we can figure out about 15 different medical conditions, and so on and so forth.

We use all that information. We recomputed the costs of their drug consumption, and we have alternative plans in each one of the markets, and we can do a little bit of switching. We have one opportunity to look at that. We also analyze what happened with this out-of-pocket spending. We apply some elasticities for the demand of drugs. We also look at the case with zero elasticity. We look at many of the issues.

So here we have some of the characteristics of a plan and the plans not included in our sample, in general, with a few exceptions, the plans are not very different from the rest, so let's call them representative.

In these simple regressions, when we look at the magnitude of the change in overspending, the alpha is the parameter we want to figure out. Here we have changes in health conditions, and that's essentially the effect.

We are getting some data around $300, so in principle just from one year to the next, the individuals are going to be reducing their out of pocket expenses by about $300.

We run the same thing for a subset of individuals with stable conditions, and what we are looking at here is individuals with a very small variation in this risk index. The last column of this implies that there are no changes in any of the ten medical conditions that we are tracking. It looks like that's not relevant, so this reduction is independent of whether we have stable conditions or not.

You're going to estimate this from the table that essentially 80 percent of the people are able to reduce their expenses going from 2006 to 2007. 20 percent go up, but this could be perhaps due to some shocks or something like that, but most of the people reduce it.

Since we have these demographics and we have the information, we want to figure out, is there any pattern? Can we figure out whether there are some individuals that are able to take advantage more than others and so on?

Out-of-pocket changes essentially vary with demographics. Here we have the common conditions like cholesterol and diabetes. They are the ones who realize larger reductions in the out of pocket expenses. If you look at Alzheimer's, they were not very far away, so they fall actually within the bulk.

We are not claiming that people with Alzheimer's or dementia are smarter than the rest, but that we have institutions who are actually helping those individuals.

This is something that happens; it's very interesting. Another important effect in the reduction happens for age groups 80 to 85 more than people of about 65 to 70 or 70 to 80. 80 to 85 are actually one of the groups that reduce the out-of-pocket expenses more than anybody else. We have evidence that this is likely because of these institutions, family, or others who help them to make those choices.

I'm sure I'm skipping things. It's impossible to go over everything here.

The next step we were going to look is how important is switching and what people are going to engage in switching? And you can see here that the top two panels are the distribution of overpayments in 2006 and 2007 for those individuals who switch, so clearly there's a switch in the distribution when we go from 2006 to 2007. The switch is not that important when we look at individuals that don't switch.

Let me summarize the results that we find. Most individuals are going to reduce their expenses by about $436 by switching plans. That doesn't mean that individuals who switch are the only ones who save. We saw individuals who did not switch plans reduce their out-of- pocket expenses by $137 on average, so sometimes they're becoming less expensive. There is another issue that we can address here: What is the effect of inertia? We could say, "Well, I signed up for this particular plan and I'm going to stay for this plan forever."

Another thing we find is if your plan goes down into the range, it becomes suddenly much more expensive, you tend to move out of that plan. At least that's what we find in this first year.

That's the effect of inertia. Health status is not important, and another interesting thing we find is that people who acquire a condition are e the ones who would use their out-of-pocket expenses more.

One way to interpret this is conditions like hypertension. You know that they are coming so those guys actually sign up in 2006 for plans that were a little more expensive than necessary, but even though they didn't get their medications for hypertension in 2006, but they got it in 2007, and at that moment their savings are much larger.

I think the thing we do is robustness for 2007, we can actually look at expected consumption, taking the consumption in 2006 and see what happens, and we changed the elasticities which is another thing by the way.

DR. SCOTT MORTON: Wonderful. Thank you.

DR. HOADLEY: So, my disclosure is I'm a political scientists here in a room full of economists, and I don't know what that's worth.

Basically what the authors of this paper start from is the premise that the Part D experience has been largely positive. I'm going to test the question of whether the trust environment is part of the positive side of that, that we're really seeing people respond well to the choice opportunity that Part D creates.

I would just like to put a few caveats on the table on challenges that are remaining, and one is that we're seeing a lot of increasing premiums for beneficiaries. We're seeing a lot of volatility for the low income, although this paper didn't really address the low income side.

My work suggested choices are still quite confusing, and that the cover gap, the doughnut hole that people face, are still huge challenges, although the health reform does address the coverage gap in a couple of these other things.

To make a few observations, just about the Part D experience from how I see it related to this question of whether choice is working, over the course of the five years in the program, we see that beneficiaries seem to stay in plans in the face of significant premium increases, despite the fact that the overall cost of the program has been on a fairly reasonable track.

The average premiums that beneficiaries face go up 44 percent in the first five years, another 9 percent projected if people don't switch for 2011. And we have individual cases of plans with as much as a couple hundred percent increase over the first couple of years, and yet retaining fair enrollment despite those very, very substantial increases.

We did focus groups with beneficiaries, and we asked them very specifically about their experiences of changing plans. Did they do research? What did they do at the end of the year? What if their premium goes up? And the consistent thing we keep hearing back from this is that that's too confusing to do the research, to even go into the question, and that they have no bias toward staying put.

It's just not about the confusion of doing the research. It's also about the comfort that we've heard some talk about in some of the other papers today, but just the comfort that the plan they're in, the intangibles that keep people in the plan they're in. They're in the AARP plan; they're in the local BlueCross/BlueShield plan; they've learned how to operate with that bureaucracy and they really don't want to switch.

We've also heard that in behavior in terms of dealing with the plans, if the things that would typically drag you to changing plans or at least drive you to try to get an exception to get your drug covered, the process of making those kind of changes is difficult for both the beneficiary and the physician, but the temptation is to change the drug rather than change the plan.

And of course that is one of the things that is going to help them look like they're financially in better shape. Sometimes those are perfectly good changes. Sometimes they may be not ideal changes.

To the point of the confusion in the plans, this was average premiums. These particular ones were unweighted. They would be similar if we used weighted by enrollment for what I consider the basic benefit plans, and the first two buys versus the enhanced plans.

The enhanced plans, by definition, are the higher value. They have to have something that makes them better than a basic plan, and yet the price for those are almost identical, just a few percentage points higher. When you go to a plan with coverage in the doughnut hole, then you will see a huge jump in the monthly premium despite the fact that it can be hard to differentiate some of these plans. I could go into the plan names.

In 2010, we have one basic plan that was called Humana Enhanced and we had plans with the designation value that are both in the lower cost plans and the higher cost plans.

This is just a point of data which shows that for enhanced plans, in most cases, you actually pay higher copays for your drugs than you do in a basic plan. Somehow this has a higher actual value, but you actually pay more out of pocket for your drugs.

This is about the gap coverage, which mostly at this point, is for generic drugs only, and all eight plans on here have gap coverage. The dark bar is sort of the share across that you face in the gap for a basket of generic drugs.

As you can see, in four of the plans up there, even though you have coverage in the gap, you're still paying 96 percent of the cost of those generic drugs in the gap, despite paying $40 a month to get coverage.

This analysis is only reported in the first two years how many people actually did switch plans. They haven't reported those numbers since then, and so far I haven't seen anybody that's actually gone into the claims data to actually calculate this. But if you exclude the low income folks, about a million people or about 7 percent switched plans after the first year, and about 6 percent after the second year.

Again, I don't know how to judge that, as a high number or a low number, but given some of the volatility in this market, the number of choices out there, and the number of people who appear to be paying too much, you might expect that this is a really vibrant market to see more switching. Although, this is may be not inconsistent with some of the electrical markets we saw in the last paper.

So the paper talked about this shift in the amount of overspending, in other words, how much more out-of-pocket costs are you paying in your current plan compared to any of the other array of plans that you could have switched into? And this overspending reduced from '06 to '07. The paper asks whether that reflects decisions to choose new plans or changes in both available plan options, and finds that the switches did a better job in reducing their overspending than the non switchers.

I find this to be a really interesting paper, and I hope it's true that we actually are seeing this kind of switching. It's something that hasn't been studied enough and it needs to be studied more, so I'm glad to see this.

My concern is with the sample that we're working with in that paper, because it's basically working with only the plans operating under one of the PBMs that operates in this market. It does include the direct plans that that PBM sponsors, like Caremark sponsors the Silver Script plans and they're not only called the Caremark plans, and there is also a claims administrator for some other set of plans.

I don't actually know which ones those are, but what you've got is a relatively small set of the overall market, and so I am wondering whether it's constrained and unrepresentative set of options in that market. I don't know if this really illustrates what's going on in the full Part D market or what's going on in this sort of sub market.

And a fair number of the plan switching that we're looking at here are switches within the plans offered by one sponsor, by Caremark. I think that's a rather different switching environment than what the overall marketing is allowing, which is a set of choices across a number of different sponsors.

Now, there's a little bit of that in there, but it's hard to know how much is in there. Between your '06 and '07, the sample changes a bit because of the new plans that that administrator takes on.

I won't go into this in any detail, but the one point I wanted to make is that the Silver Script, the basic plan in '06 versus '07 was actually an exception to the overall rule, and the plan went down by about a dollar in premium. And their enhanced plan is actually one that went down by $20 in premium.

So what you've got is a typical context of choice demand here. Silver Script added a third plan that year, which was a lower robust enhanced plan, and had it at a cheaper price than their enhanced plan had been the year before. And we could go through the rest of the details on this.

Again, the sample does have a few other samples, but since I didn't know which ones those were, I couldn't compare those. So, I think, as provocative as this paper is, it really does try to answer an important question from the Medicare perspective and obviously with broader implications for market behaviors, but I think the paper does try to start doing this.

How do you sort out what's going on with this improvement that they sell? How much of this was a conscious decision to research and switch plans? How much of that was a choice within sponsors versus across sponsors and how should we be generalizing from those things? How much of this is about drug use changes that people may be making across the two-year period relative to formularies?

That's potentially a good choice to make if they're making it in consultation with a doctor and going to a perfectly good substitute drug. It could be a not ideal choice for some other instances where you're bringing your numbers more in line in terms of this overspending.

And how are the results in plans changing? Some of those premium differences I showed you on the previous screen are some of the cost sharing differences that you're passively accepting because your plan is making choices. The theories put that question as one of the things out there.

But before we draw as many conclusions about what's going to happen in this market, I think we really need to understand that. Secondly I think we need to think about whether these results are unique to the '06-07 period.

There was a lot of shake out in the market. A lot of plans like this one came in with not very well designed premiums, perhaps not well defined benefits overall, and were making adjustments from year one to year two. They are continuing to make those adjustments every year, but knowing there's this kind of separation in the first two years.

And I'm wondering how the same kind of analysis would look if we were looking at '06 versus '08 or '09 or even just the years from '06 to '08 or

'08 to '09. Finally, the results would vary if we looked at a variable set of plans or plans within different sponsors.

I'll stop it there.

DR. SCOTT MORTON: Thank you very much. Do the authors want to make any comments about their discussants' response or things they omitted to say when they this the microphone? OK. Everybody is happy.

(Discussion off the record.)

(Applause.)

(Whereupon, at 5:26 p.m. the conference was adjourned.)

CERTIFICATE OF REPORTER

DOCKET/FILE NUMBER: P085800
CASE TITLE: MICROECONOMICS CONFERENCE CONFERENCE
DATE: NOVEMBER 18, 2010

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the steno notes transcribed by me on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: DECEMBER 6, 2010

____

DEBRA L. MAHEUX

CERTIFICATION OF PROOFREADER

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

____

ELIZABETH M. FARRELL