FEDERAL TRADE COMMISSION

FIRST ANNUAL

FEDERAL TRADE COMMISSION & NORTHWESTERN UNIVERSITY

MICROECONOMICS CONFERENCE

Thursday, November 6, 2008

9:00 a.m.

Federal Trade Commission

FTC Conference Center

601 New Jersey Avenue, N.W.

Washington, D.C.

FEDERAL TRADE COMMISSION

I N D E X

P R O C E E D I N G S

\- - - - -

**WELCOME AND OPENING REMARKS**

MR. BAYE: Well, why don't we go ahead and get started. My name is Mike Baye. I'm the Director of the Bureau of Economics here at the Federal Trade Commission, and it's an absolute delight to be here to kick off the first annual FTC-Northwestern Microeconomics Conference. As you can see, we strategically put the firm "first annual" in the title as a commitment mechanism so that this will be ongoing even after I leave in a month and a half. Hopefully, that commitment mechanism will work.

Just a couple of announcements. I, first of all, want to, on behalf of the Federal Trade Commission, thank Northwestern University for their partnership in this ongoing endeavor, and in particular, to thank the Searle Center and Henry Butler for -- where is Henry? Is he here somewhere? There you are, Henry. Had breakfast with Henry this morning. Thank the Searle Center for their support in this ongoing relationship, and, also, the Center for the Study of Industrial Organization, in particular, Bill Rogerson, who was instrumental in helping forge this partnership with Northwestern University, along with Aviv Nevo and Scott

Stern. So, we really are thankful for those partnerships.

I know Chairman Kovacic does as well. I believe that knowledge is really the key to good decision-making, and I do believe that it's important for government and some of the best thinkers in the world to get together to analyze and to think about some of the important issues that arise in economics. And I'm certain there will be tremendous fruits of the research and dialogue that goes on here today and in the future.

I'd also like to extend my gratitude to the scientific panel, who played an integral role in helping put the program together, in selecting papers and so forth. Let me just briefly indicate who they are: Susan Athey, Pat Bajari, John List, Carl Shapiro, and Scott Stern, again, people providing public goods for the benefit of science. I'm really thankful for that.

And then, just lastly, let me thank the people here at the Federal Trade Commission and in the Bureau of Economics that put a lot of time in to make sure that this event actually was kicked off and was more than an idea. Special thanks to Chris Adams, who played an instrumental role in doing all the organizational work, putting together scientific committees, and so forth.

Dan O'Brien, Matt Weinberg, Rob Leitzler, Loren Smith, Marissa Crawford, Mary Villaflor, Neal Reed, Matt Eaton, Tammy John, and Alethea Fields, all played an important role in making us comfortable and putting together a great program.

And it's my distinct pleasure to be able to introduce Bill Kovacic, our Chairman, who's going to be kicking off this event for us.  Chairman Kovacic is both a gentleman and a scholar.  He has served in various capacities at the Federal Trade Commission.  He's served as a staff attorney; served as the General Counsel to the Federal Trade Commission; he served as a Commissioner; and most recently, he's serving as the Chairman of the Federal Trade Commission.

And I know of no better person to kick off a microeconomics conference than Chairman Kovacic.  He has a true love for research.  He has a true love for knowledge.  And he's an academic in the very best sense of the word.  He's a distinguished attorney, as you all know, but what you may not know is that he's co-authored with a number of illustrious economists, including Patrick Ray, Bob Marshall, and Leslie Marx.  He hasn't offered to co-author a paper with Mike Baye, but despite that, I will introduce Bill, my friend and our Chairman.

CHAIRMAN KOVACIC:  Thank you.

(Applause.)

CHAIRMAN KOVACIC: I'm also enormously grateful to Northwestern and Henry and his colleagues for assisting us in putting the program together, certainly to the scientific committee that has assembled a fabulous agenda, and to our Bureau of Economics, and most of all, to Mike Baye. Mike is the very latest worthy successor in a tradition that's brought an illustrious collection of civil servant scholars, academics, to the position of the head of the Bureau of Economics, and it's no accident that an event of this kind would take place with Mike's guidance, and I do, indeed, expect it's the first of many great events in the future.

A couple of thoughts about the motivation for putting this together, at least from my own point of view. I think there's been a growing realization -- and maybe it's a consequence of having academics who have served as Chair here, certainly Bob Pitofsky, Tim Muris, and myself -- that if the Commission was to do effective work in a number of particularly difficult policy areas involving consumer protection and competition policy, especially in areas involving enormous technological dynamism and organizational change, that it would have to increase investments in building knowledge; that when you looked at the budget of activities, that the element of production or consumption that's involved in the prosecution of cases could only take place sensibly if we were making capital investments and making more of them; making investments that in any one budget period would have a long life, especially those associated with building knowledge.

And I think there's been, over the past 15 years or so, a very healthy norm or custom developed inside the agency to increase awareness of that, and this program is a manifestation of that awareness.  Why do this?  First, I think it's necessary to our capacity to deliver good programs, be they in the form of reports, enforcement programs, and competition or consumer protection, in advocacy before our legislature or state bodies or, indeed, to have influence in a larger global setting of shared authority; that without major continuing, substantial investments in building knowledge, that we won't be able to do good work in this area.

The last element of that that I mentioned becomes increasingly important.  Regulatory authority in the United States and around the world is highly decentralized.  We share authority for what we do in both areas of our competence, competition and consumer protection, with an astonishing collection of other federal, state, and local institutions that have concurrent authority, with no trumping mechanism that dictates that any single institution controls the decisions of the others. Internationally, it's simply the same setting.

And in all of these areas, especially in a world of over 100 competition authorities now, more consumer protection authorities, you don't exercise influence by compulsion. You exercise influence by persuasion. And I'm convinced that the competition or consumer protection agency of the future that aspires to really shape what other organizations do will be the one that makes the major investment in ideas. It becomes the focal point for doing good research, issuing good reports, and demonstrating, through its selection of programs, that it does have intellectual leadership.

And my own aspiration is that in the whole field of competition and consumer protection groups, that the Federal Trade Commission will be paramount in that effort. Where do we stand now in that effort? I think of a story told about a journalist who interviewed the curator of The Hermitage Museum in St. Petersburg and asked the curator, said, "Is your art collection the greatest in the world?" And the curator says, "Well, that's a difficult question. I'm really not sure, but I can assure you of one thing. We are not the second." And I think in our own, as a public institution, we are not the second, but making investments that develop true preeminence is going to be the way that you exercise influence. In other words, you can't shout people into agreeing with you, but you can persuade them. This kind of activity is part of that process.

How do we mean to do this? Through our own research, through a number of the researchers that we have here, folks well known to you, like Pauline Ippolito, folks like Jim Lacko, Jan Pappalardo, any number of others whose work shapes the way people think about the field, through workshops, but importantly -- and that's part of this program -- through links to other research institutions. The relationship with Northwestern is what I hope will be the first of a number of links that we develop with major research institutions.

When you compare us to other jurisdictions around the world, the United States competition and consumer protection community enjoy a major advantage, and that is North America is absolutely unsurpassed as a focal point for research involving industrial organization, information economics, other foundations for what we do. And to some extent, we have drawn upon that knowledge in a variety of ways, but I think by achieving deeper integration with the research community in North America, a two-way exchange of ideas, talking more about what we do, drawing in more in real time what researchers are doing, working through recruitment and the attraction of the best graduate students into our programs, we take a major step ahead in developing the foundation for establishing true intellectual leadership.

So, the program that Mike, Northwestern, and their colleagues have set in motion today I see as being an absolutely crucial, valuable part of an effort that I think will have a very long life to build true intellectual leadership and ensure that our programs, if they falter, do not falter because of a lack of effort to build a good base of knowledge. I look forward to being able to sit in on some of the sessions today and tomorrow, and I'm enormously grateful to the researchers who have come to present their work here today.

And thank you, in particular, as Mike was saying before, making the investment in the common pool of knowledge that we'll all draw upon, but in your efforts to assist us in doing what seldom happens in this city, which is making long-term capital investments that will serve the institution well over a long period of time, resisting the impulse simply to make investments in activities that yield immediate, appropriable returns; in other words, to build a foundation that will last for a very long time for the benefit not simply of this institution, but the citizens and consumers we serve.

Thank you again, and thanks to Mike and to the entire team.

(Applause.)

MR. BAYE: Thanks, Chairman.

Our introductory session will be held by one of our partners, Scott Stern, who will be talking about the market for ideas. So, Scott.

**INTRODUCTION**

MR. STERN: So, thank you very much for the opportunity to talk for a brief amount of time today, and I just want to kind of echo both the comments of Mike and Bill regarding this effort, and I think that sort of this is a very exciting beginning for a really interesting foundation.

I'm going to give a paper today that is a little -- I am going to be very up front. This paper is quite speculative, but I think very interesting, and I think potentially important for this audience and for the area of thinking about competition policy and even some of the consumer protection issues that arise in a world where an increasing share of the economy is grounded in the production of knowledge-intensive goods and services that depend itself on really the development of the incentives and diffusion and commercialization of new ideas.

So, I'm going to start with some very basic ideas here, which is that we sort of know kind of back when we were taking econ undergrad that in principle, markets for ideas or technology have, in principle, potentially very high social returns. In particular, given the uncertainty associated with the innovation process, a single idea may be valuable to many users and also in many applications, often in contacts that are far removed from the locus of invention. So, given the uncertainty of innovation, you might come up with an idea in context A, but what we know from studies of innovation is very often the highest benefit is realized by playing that in a very different context.

In particular, the value of that idea depends on somehow matching it effectively with complementary assets, and moreover, if somehow people who are developing ideas that ultimately were applied and there was the right price for that idea, that would provide very effective signals for future investment in idea production itself. So, that's sort of the very principle, kind of high social return activity.

At the same time, markets for ideas are actually, relative to almost every other transaction you can think of, pretty darn rare. They are not absent, and in areas like I've studied in biotechnology, they are actually quite common, but outside of a few very narrow segments of the economy, very, very little transaction in ideas.

More importantly -- and that's going to be the main point of what I'm going to talk about today in the brief amount of time -- is that most of that trade occurs not really in a market, you know, so we talk about markets for ideas, but what we really mean is I'm some little entrepreneur, and I finally found someone to buy my new idea, and I'm pretty much -- it's a bilateral transaction, and very little of our analysis has really distinguished between what it would mean to organize knowledge, exchange, and diffusion as a market as opposed to a series of isolated bilateral transactions.

Intriguingly -- and I'll come back to this, once again, in the small amount of time -- the most robust way we know how to do this -- and it was already alluded to in both of the earlier comments -- is something that we most -- essentially everyone here participates in: The republic of science. Interesting point about that market: The price of the ideas is exactly equal to zero.

So, what we're going to do here -- this is, by the way, joint work with Joshua Gans from Melbourne Business School. We combine two distinct literatures. On the one hand, this paper was motivated -- and I'll be very explicit about that -- that I had the opportunity to sit through a talk by Al Roth on sort of frontiers of market design, and I heard that talk four times in the course of a year, and I'm very slow, but by the fourth time, I figured, huh, that's something that people who think about innovation might think about, and I'll kind of come back to that theme.

And I'm going to sort of take the ideas around kind of economic analysis of the requirements and challenges of market design, which is something that I think a lot of you here will have been familiar with at some level with our understanding of markets for technology. And we are going to come up with kind of three propositions that I'm, in the short amount of time, not going to be able to kind of really develop each of them in their full development.

But the first is that the very nature of what ideas are undermines the market for ideas, and that's an important point, that there's a kind of fundamental source of the ability to allocate exchange in a market with multiple players on both sides of the market.

The second is that the most robust market for ideas are those where the ideas are free, and that's going to raise this notion that Roth brings up of what's called repugnance.

And then the third thing is that formal intellectual property rights may not simply facilitate isolated transactions, which is sort of how a lot of work in economics, I think, is going to shape the agenda there, but they actually play a crucial role in overcoming challenges to establishing a market where essentially the outside options for both buyers and sellers is potentially other transactions in the market, okay, and I'll come back to that.

Key point, you know, sort of for this audience is how does -- so, what I want to do is sort of ask the broad question and really, if all I get out of this, you know, kind of short amount of time is to raise the question up for further analysis, is say, how does market failure in innovation markets impact effective competition policy towards innovation and technology?

Let me just kind of -- to give you the context here, in the licensing of -- the guidelines for licensing of IP, and that's specific to IP, there's no specific document related to kind of knowledge and ideas more generally, intellectual property has important characteristics, such as ease of misappropriation, that distinguish it from other forms of properties. Those characteristics can be taken into account by standard antitrust analysis, however, and do not require the application of fundamentally different principles. And I agree with that statement, in general, but I would argue that I think the competition policy and antitrust community has, by and large, abstracted away from some of the kind of thorniest issues in how we develop, diffuse, and commercialize ideas when thinking about antitrust applications towards innovation markets.

Okay, so let me just take two seconds on market design. I am going to kind of go through this relatively quickly to kind of develop the idea. So, as some of you are aware in this room, some of you have been actively involved in this, as I look over at no time table over there, mechanism design offers powerful abstract insights into the efficiency and limitations of alternative market allocation mechanisms when buyers and sellers possess private information.

And as most people here know, a lot of game theorists have gotten pretty darn involved in market design, from FCC spectrum options to the work by Al Roth and others on things like the National Resident Matching Program, to even the plethora of economists we've seen not entering the civil sector, but, in fact, going into the private sector or having relationships with private sector firms to develop markets or auctions, for example, for Internet advertising positions, such as Hal Varian and Susan Athey.

Now, rather than simply apply the theory, okay, real world applications have opened up, I think, new insights into the requirements for efficient market operation. In other words, how do you create a market in which both buyers and sellers essentially have endogenous outside options and raise new theoretical challenges? And this is an example that Al Roth has developed in some detail -- let me just -- has developed in some detail.

What's the problem in kidney exchange? First, the prices are zero, by law. If you need a new kidney and somebody's willing to give it to you, you have a very high chance that their blood type and other characteristics are incompatible with your own body, and in principle, then, and if there are two people who have that, you, in principle, could have a very good market exchange, but you have to sort of organize the market effectively. So, many kidney diseases have willing donors, but there's limitations on donations due to incompatibilities.

The market for voluntary living donor exchange, which they've set up, has dramatically, in a relatively short time, facilitated the ability for people to get new kidneys, and, you know, live. So, this is a life-and-death issue for 10,000 people. We could have 10,000 people die per year because they don't make it off the kidney donor waiting list. You do a little bit of market design, that resolves the fact that incompatible types can lead to very high exchanges for trade across people, and even though there's no prices, what happens is Michael and I will make an agreement that your donor will give to me, my donor will give to you, we do the operation at exactly the same time, and we achieve very high returns.

So, Al Roth in the HAM lecture last year developed sort of three criteria for effective market allocation system, and I think if you haven't read this article, it's actually quite effective, I think, in synthesizing a lot of the work that's been done in market design over the last ten years. One is, what do markets need? They need thickness -- that's something I think we knew probably already; we need lack of congestion, essentially, individual transactions have to be set up so they have enough time to look for an alternative offer. Exploding offers are disasters from the viewpoint of social efficiency. And finally, market safety, which is a simple way of you have to be willing to basically report your type, okay?

And finally, one thing that Roth and I think other people have sort of kind of taken away is that an important lesson for many real world market design problems is that there seem to be important constraints on these markets grounded in social behavior, the idea that he talks about is repugnance, that social norms play very significant informal and formal restrictions on the ability to use prices to facilitate allocation. The simplest point of this, on the one hand, when two people get married at a price of zero, we all think that's great, but the market for sex is mostly prohibited. I haven't looked at the California -- San Francisco proposition on that point, but nonetheless, okay.

So -- okay, I don't have a lot of time. What I'm going to do is the following. I've sort of misallocated how much -- okay.

So, what I want to do is say, on the one hand, what do we mean by a market for ideas? It's going to be characterized by -- once again, by this distinguishing feature between isolated transactions with may be very, very high search costs, okay, but just isolated transactions so that if you -- if failure in bargaining occurs between the two members, if the idea is really useful, the alternative option is seeking, for example, an alternative buyer, versus -- excuse me, bilateral transactions versus a market where the option is endogenous.

And so there are going to be two features -- and once again, we could go through more of these, but I just want to highlight two -- there are two lots of ideas which impact the challenge of market design in the market for ideas.  The first is something we call sort of value rival, kind of coming up from -- kind of thinking about ideas around nonrivalrous -- nonrivalry of ideas.

So, in other words, in biotechnology, even if I have a pretty strong piece of intellectual property, but there's some tacit knowledge around it that I would have to sort of disclose, one problem I face is that if I approach a pharmaceutical company and start telling them about my idea, not only, right, what I would really like to do is approach actually many pharmaceutical companies all at the same time, but the value of each of those potential buyers from buying my patent is declining if the general knowledge that's associated with that idea is also being diffused through the bargaining process to my -- to the buyer's potential competitors, right?

So, in other words, if I review -- right?  So, if I have a secret and I want to share it with Carl, and Carl's competitor is Mike, right, is Michael, and if I also say, listen, I have Michael is also willing to buy the idea, Carl's like, well, now the secret's gone, and so I don't even want to transact with you anyway.  And what that does is mean that the very fact that -- so, the misappropriability problem actually degrades the bargaining process.

A second problem, which we all know, any of those who have teenagers who do Napster, and user reproducibility, right? The cost of -- there are certain ideas where the cost to a buyer of reproducing the idea effectively as a seller is extremely low. So, it's not just that the marginal cost to the initial producer is low; it's that once you buy the idea, you can kind of walk up and say, I used to be a buyer, and now I'm a seller, and that's going to reduce, and essentially that means that's going to limit the scope for transactions in general, okay?

Okay, so I'm going to sort of talk about this. So, what's going to happen here? The value of ideas -- so, I talked a little -- okay. So, you know, so a fairly highly cited article from just this year in the kind of law literature, Mark Lemley and Nathan Myhrvold talk about what's the problem in developing the market for intellectual property, and they say, imagine a stock market in which buyers and sellers couldn't find the prices out at which anyone sold a share of stock. Well, that's actually -- since that was written, that's probably more true these days, right? If you wanted to buy a share of stock, you would have to guess what it was worth. Willing buyers and sellers often miss each other. The prices vary dramatically from sale to sale, and private -- people with private information can exploit others, and, of course, everyone uses that. Surely no one would intentionally design a system in which trades had to be kind of blind in that way, right, which is a lot of uncertainty and inefficiency.

As best I can tell in my own research, and I think probably anyone who's sort of looked at this in any detail, as they say, patents exist exactly in that market. If you talk to 20 biotech companies about what they went through to get their license -- and that's probably one of the more developed exchanges -- you essentially have wide variation in exactly what type of terms were put on them and exactly what kind of baseline pricing there was, right?

Second, ideas can be reproduced by users at essentially zero marginal cost. There are going to be significant limitations on whether the seller can control how others use or distribute the idea, right? So, what's going to happen here, if I know that after I sell it to the first person at any price -- this is sort of the work by David Levine, is sort of this trade -- because you sell it to the first guy for a high price, and then he sort of maybe sells it off to others with some imperfections, but the key problem is if I sell it to the first buyer, that buyer becomes a seller, and the marginal cost equals zero, and we're in Bertrand. That, of course, inhibits my incentives to sell in the first place.

So, you can sort of unravel the entire market if the potential seller -- potentially buyers can also become sellers and drive the price to zero. User reproducibility in some sense results in a failure of what Roth would call market safety. Individuals have incentives to engage in strategic behavior that undermines the welfare arising from allocation, and, you know, you look at what happened in digital music in the late nineties, and even today, that's what's happened, right? All sorts of problems on the insider side, because every single buyer can effectively become a seller.

So, very quickly, we do know that there are some people who think -- you know, initiatives that people are taking to do that, from normal intellectual market exchanges to something that people really haven't looked at, for example, the competition policy of very organized exchanges in very key trade conferences. Most of the transactions in technology that we think about occur in the context of formalized -- or occur against the background of industry associations that really have mostly escaped regulation in terms of their economic function. Transparent platforms, to a certain degree, standard-setting organizations, which is an area of very recent interest, okay?

Let me just go -- let me, in the interest of time, let me just finish out in the last two minutes here on just them pushing this idea around free. So, that's Benjamin Franklin. Certainly anything that Benjamin Franklin says is probably at least well said. "As we enjoy great advantages from the invention of others, we should be glad of an opportunity to serve others by any invention of ours, and this we should do freely and generously, right?" Ben Franklin was the ultimate enlightenment thinker, right? And that's -- you know, 200 years later, we now say with fewer words but less elegance, "Information wants to be free."

Now, the idea that buyers of information, right, the fact that users would like information to be free is not that surprising. They want a low price. The fact that many of the most strong advocates for, for example, for open star software are the producers is something that as economists we have to think about in terms of their incentives.

Not only that, but not only do they think that ideas should be free, but there's a huge discontinuity between zero and any price, okay? So, for example, Chris Anderson in The Long Tale says, right, charges -- you know, essentially, in many cases, there's a -- the psychology of free is powerful. The truth is zero is one market, and any price is another. And it's true. Micropayments are almost a complete failure. What you see is people either have zero on their idea or they go out, get a patent, engage in very big-time, you know, kind of thinking about it, and sell out for a very big price, but kind of the kind of intermediate range of idea exchange is essentially missing. That's a missing market in almost every context I can think of.

And the question is why? Is it something -- right? And what I just want to kind of in the -- and I know I've gone over my time a tiny bit, but what I want to do is just kind of raise up an idea that Roth first introduces in the context of thinking about things like kidney exchange, markets for throwing dwarfs, markets for, you know, all sorts of things, is that there seems to be a part of ideas where you can sell for free in which people have -- in which there seems to be markets, but the prices are free. And so just -- and it kind of raises up this notion of what he calls a repugnant good.

So, let's try this. So, this I'll end on. So, should the following -- just in your own mind, should the following activities be permitted, in general? So, Steve Jobs, we know, charges a price much greater than marginal cost for the iPhone. I imagine most people here are comfortable with that. How about a pharmaceutical firm charging a price much greater than marginal cost for a malaria treatment that was completely discovered with public funds? Third, how about the right of a record label to prohibit artists from playing their own music with heavy penalties for infringement?

How about licenses for university-developed, sort of scientific-developed -- science-developed, general-purpose research tools which involve very significant -- where the form of the contract that's agreed upon involves very significant restrictions on the ability to publish follow-on scientific research? There's a well-known case regarding the INCA mouse that deals with that.

How about an auction between you and your health insurance company to have exclusive access, either you or the insurance company gets access to your genetic profile? How about secret payments by the Government to journalists or bloggers to express particular opinions as their own? How about the sale of credit for a discovery by a student to a faculty member?

All those are markets that somehow I imagine that at least some of you might have some problems with what's occurring at the bottom. And I would just entertain that why we have problems with repugnance in some markets and not others, all of which have to do with the production and distribution of ideas, is an interesting area going forward.

Thank you very much.

(Applause.)

**PANEL SESSION ONE: ESTIMATING DEMAND AND PRICING STRATEGIES**

MR. BAJARI: Good morning, everybody. I'm Pat Bajari from the University of Minnesota.

UNKNOWN SPEAKER: Is your microphone on, Pat?

MR. BAJARI: Can you hear me now? Okay, great.

Pat Bajari from the University of Minnesota, and we're going to have next a short panel discussion on merger simulations, and I want to just first briefly introduce our three participants. The first, to my left, is Mike Vita, who's an Assistant Director for Antitrust at the U.S. Federal Trade Commission in the Bureau of Economics. Mike has published numerous academic papers within industrial organization and antitrust. In particular, he's supervised a number of merger investigations in which merger simulations have been used. These include pet foods, ice cream, spices and hospitals.

Next to Mike is Aviv Nevo, who's a Professor of Economics and Marketing at Northwestern University. Aviv has published widely on differentiated product demand estimation and on merger simulations, and he's on the editorial boards of a number of leading journals. In addition, he's worked on several real world merger simulations as an expert.

And finally, Gail Slater is a staff attorney at the Federal Trade Commission. She's been at the Federal Trade Commission since 2004, and she's worked on a number of merger and nonmerger cases, most recently the Whole Foods case.

So, what we're going to do is I'm going to give everybody five minutes to make a brief statement on some of their views about merger simulations, give them a little chance to respond to each other, and then I'm hoping you, the members of the audience -- I know there's some people with opinions about merger simulations here -- would be helpful and chip in by asking some questions of our panelists. I think this will be a fun topic to discuss where people have some different opinions. So, please be thinking about this in the background and help us out by participating, because I'm sure your questions will be a lot better than mine.

MR. VITA: I'm going to do a little PowerPoint.

MR. BAJARI: Okay, great.

MR. VITA: Okay. Let me start off with the obligatory disclaimer that everybody here at the FTC has to give. These are my views and not those of the FTC or any Commissioner, and that's almost always true any time I speak.

So, yeah, as Pat said, you know, I'm a manager here in the Bureau of Economics. My job is to manage antitrust investigations from the Bureau of Economics' perspective. Merger simulation has been a big part of what we do here in the Bureau of Economics for about the last ten years or so, when some of these technologies, if you want to call them that, first appeared on the scene.

Back when those first papers were being written back in the late nineties by people like Greg Werden, Luke Froeb, at the Department of Justice and the FTC, respectively, and Jerry Hausman and others, I think people had a great deal of optimism about how much this could add to our analysis of mergers, at least certain kinds of mergers, and, in fact, Greg and Luke have a paper -- a couple papers entitled, "Merger Simulation as an Alternative to Structural Merger Policy," and by "structural merger policy," I think they mean sort of the traditional antitrust analysis where -- whereby it's, you know, centered on document reading, interviews, depositions, that sort of thing, calculation of market shares, and everything that's in the 1992 Merger Guidelines.

Now, the typical simulation exercise, you know, people who work at the agencies, you know, know what I'm talking about, but for those of you who aren't real familiar with it, what I mean, and I think what most of us mean when we talk about merger simulation, is, you know, assuming a particular functional form for demands for the products in question, assuming a particular form of competition, usually Bertrand competition, and estimating the parameters of those demand functions and then, combining that with the assumption about the nature of market competition, predicting what the price would be, and the output would be in the post-merger equilibrium.

So, this has been going on for a long time. As Pat said, I've done a lot of cases, you know, in addition to the case -- the industries Pat mentioned, you know, I went back looking through some of my notes for other cases. We've done it in cigarettes, breath mints, alcoholic spirits, in addition to all the other ones he did. And you would think that -- you know, in my shop, you can probably tell we do a lot of consumer products cases, and this is an environment, I think, that most economists would predict would be ideally suited to this kind of exercise.

So, one of the questions, now that we've been doing this for about a decade, that we frequently get from, like, our colleagues in the Bureau of Competition as well as from Commissioners is, how well does this technology work? Does it do a serviceable job of predicting what post-merger prices and outputs would be like? And only now are we really getting to a point where people are starting to address that question and can say something interesting and important about that. And the evidence on that question is fragmentary, but there are a couple of papers -- one has been published, one I think probably will be published in the next year -- that get to that issue, and frankly, I think the -- you know, the results of that research so far, it's a little disquieting to those of us who have been using this method and have been recommending its use.

The first paper is by -- was published by Craig Peters of the Department of Justice, who went back and looked at a number of consummated airline mergers and went ahead and -- what he did is he took data from the premerger world and went through the simulation exercise, estimated demand functions, and then simulated the post-merger environment, and he did it under -- you know, the details of exactly what he did aren't terribly important, with the possible exception of he assumed a static Bertrand codec, which I'll get back to real quickly when we talk about his results.

So, he estimated -- you know, he estimated the demand functions, simulated the equilibrium in the premerger and post-merger world, and predicted the prices. Then he went ahead and compared the predicted prices from that modeling exercise to the actual changes, and he found -- and here's a table that I reproduced from his paper -- there's -- frankly, you know, it doesn't appear to have matched up all that way.

You can see the first line, the logit models, the first -- the second column, where it's labeled "Observed," those are the actual prices that, you know, actually obtained from -- you know, from looking at post-merger data. The logit and GEB, those are simulated prices based on a couple different modeling assumptions. And you can see, if you scan that, it's -- you know, it's some pretty big divergences between the observed and the actual.

The second paper, the second piece of evidence on this subject is being done here at the FTC by Matt Weinberg and Don Hosken, and Matt's going to be presenting tomorrow where he's going to talk about his work in a lot more detail, so I'll just talk about it real quickly. We looked at -- Matt and Dan looked at two consummated mergers in consumer products. One is motor oil, from the Pennzoil-Quaker State merger, and the second one is maple syrup, from the Aurora-Log Cabin transaction. Both these deals went ahead and were consummated with no enforcement action.

Again, you know, in their paper, they used the standard sort of thing. They estimate demand functions under a couple different functional form assumptions and go ahead and predict the post-merger price, then go ahead and calculate or estimate the post-merger price using private label products as the control. And what they find in the results is that, you know, again, they get a couple -- a couple of the predictions seem to be, you know, pretty close to what actually happened, but more generally, and the bigger problem is, the actual price change for oil seemed to be pretty large and pretty small for syrup. The simulated price changes got those things exactly reversed. So, I mean, that's -- again, I -- that's somewhat troubling. You would hope that any -- you know, you don't -- any serviceable prediction tool would at least get the rank ordering right.

I'll just -- I will go through this real quickly. I mean, Peters in his paper does -- goes through -- you know, tries to figure out exactly why, you know, the observed prices didn't match up all that well with the predicted price changes, and he does a really nice exposition of that. I'll just skip through that.

His conclusion is that in large part, though, the inaccurate prediction may reflect the fact that the premerger firm conduct wasn't Bertrand, which is our conventional assumption in these kinds of exercises, and there may have been something like tacit coordination going on. Matt Weinberg and Dan Hosken, in their work, they don't think it's -- you know, one of the possibilities is cost or demand might have changed. In their paper, they don't find evidence for that, and I'm not -- Matt can talk more tomorrow about what he thinks was really going on.

So, the bottom line, I guess, you know, as we continue with this -- you know, with this process is, we'll continue to do merger simulations in the bureau of of economics whenever we think it's, you know, appropriate and possible and the data permit it, but, you know, has it fulfilled the promise that it -- you know, some of the innovators predicted ten years ago where it could replace or substantially replace conventional analysis? I don't think so. It's helpful, it's a useful piece of information, but we're not really to the point where I think we can tell people, yeah, you can really rely on this as a fairly accurate predictor of what's going to happen in a post-merger environment.

So, that's probably about my five minutes. So, I'll take a seat.

MR. NEVO: Okay, so I guess I'm -- I'm on this panel, I guess, supposedly to be the big defender of merger simulation, and I might be or might not be. I'm not sure. I mean, I haven't really made up my mind yet. So, we'll see how it sort of evolves.

I guess my main point has to do with, you know, what do we really think about merger simulation? I actually noticed that we -- both when Pat introduced the panel and, you know, when we were asked to sit on it, we discussed about, you know, merger simulation, but when you actually look at the program, it talks about demand estimation and the -- something of mergers -- yeah, demand estimation for merger cases, and I think that sort of reflects a little bit sort of differing views, sort of -- for me, merger simulation is the idea that you're trying to predict what the effect of the merger will be, and I may be kind of just taking too much of a dictionary sort of -- you know, trying to interpret what the words say.

So, for me, you know, if you're doing kind of a so-called structural analysis, you know, basically Hirfendahl's and stuff like that, that's a merger simulation. You're saying if a merger falls in a particular range, you know, for Hirfendahl's, whatever your cut-offs are, then the likely effects are going to be high or the likely damage to consumers are going to be high. If you are doing sort of a so-called Staples-type analysis, okay, kind of -- as it's been called, you're trying to predict what the effects of the merger will be.

Now, then there's the narrow definition of the merger simulation, which is, you know, the one that was in the title and the one that I think Mike has already referred to, which is this, you know, specific -- you estimate demand, you take a Bertrand sort of assumption, and you use that to sort of predict what the effect would be. So, this latter one, I'm not going to be here to sort of stand and defend.

I can tell you what my thoughts about it are, but I'm not going to be defending that. I think the broader view is sort of to understand that, you know, we do need some sort of a model to predict what the effect of the merger will be. We're trying to predict something that we don't see in the world, and I think the key is to bring sort of the best economics we can to the problem, and sometimes, it might be estimating demand and putting a Bertrand assumption, and in some cases, a Bertrand assumption would be terrible. And I think, you know, we have to sort of focus a little bit at kind of understanding when, you know, these assumptions will be good and when will they be bad. So, that's just sort of in terms of kind of as a grand overview.

Let me just say a little bit something about retrospective study. So, obviously these need to be done. I mean, they are kind of long overdue, and we have to look at them. There's a bit of a problem to looking at the evidence for particular -- particular reasons. One is we forget, again, in this sort of grand view, that we're taking -- we're picking one particular method, but we ask, okay, what's the alternative? So, yes, you know, there was a table there that this type of merger simulation doesn't work well. Well, what happens if we went based on Hirfendahls? Would we do any better?

Now, it's a much harder sort of -- it's a little bit like, you know, trying to pick up an olive with chop sticks, right? I mean, it's very slippery if you're trying to get -- you know, trying to get an exact sort of something to beat up on when you just have this broad thing of, well, you know, if the cut-off is 1500, then everything about it is sort of fine. We never actually put sort of something, you know, an all prediction that someone can test later, but it's not clear that, you know, the sort of alternative would sort of do any better. The same, I think, for Staples-type analysis. You could say, well, okay, what is your prediction of the likely effects? How would that -- you know, what would that have sort of worked out to? Of course, maybe we don't have enough evidence, but we have to remember sort of in that context.

The other thing that I find in some of this discussion is, you know, I don't -- it's nice to know that we get the right effect or not. I can see why the FTC would care about that. But as an academic, I actually would like to see more and sort of see where do we need to improve our models? I think there's been too much focus on did we get the right demand. I actually think we're doing okay there. I mean, are we getting any particular cross price-elasticity right? Probably not. Now, it might be that for a lot of mergers, that's going to be sort of the key, that particular parameter, but I think where we're missing is sort of another dimension.

So, if you actually look at some of the -- I mean, you kind of mentioned them briefly, that sometimes we're not even getting kind of the relative increases right, and usually if you think hard, that's actually not about the demand. It's telling us something about what we're plugging that in, what's the supply equation. I guess we could talk more about that tomorrow when we see kind of the specific simulation of mergers, but, you know, if you look at sort of the Ashenfelter-Hosken paper, I mean, I know I looked at some of their cereal results, and on average, actually, merger simulation -- by merger simulation here, I mean kind of a very narrow view of it -- actually get things right. Where it misses is the relative, sort of the effect on the different product, and I think that's something that's not driven by the demand analysis, but actually something sort of deeper.

The other thing we also have to sort of realize when looking at these retrospectives is that we're not going to get any particular merger right. I mean, it's just asking too much. We don't have enough data. We don't know enough. There's just no way we can do that. So, we have to go back and ask, are we getting a systematic bias? So, of course, sometimes we'll be over and sometimes we'll be under. If on average we're getting these things right, I think that's the most we can hope for, and that's sort of the question that we have to ask. Are we getting this right, on average?

Now, I know that any particular merger, you know, that might not help, but, you know, that's, I think, something that we have to ask ourselves.

Okay, I'll stop here.

MR. VITA: Thank you.

MS. SLATER: Good morning, everyone. So, I'll piggy-back Mike's disclaimer. I don't speak for the Commission either, and I'd add another -- oh, I beg your pardon. I'm a low talker, so I'll start again.

I was just saying I was going to piggy-back Mike Vita's disclaimer that I don't speak for the Commission either, and I'd add to his disclaimer another disclaimer, I'm not a Ph.D. economist. I'm here as a humble staff attorney in the Bureau of Competition. So, my perspective, obviously, is a legal practitioner's perspective, and I was asked to -- you know, to think about how I see merger simulation in the legal process that I work in day to day, and that process being obviously the merger review process here at the FTC and the occasions when we go to court.

So, as I see it, there are three channels in the legal process. We have our investigations. The vast majority of cases end in either a consent or a closed decision. There are cases where we do go to court. We have the PI process. And now, increasingly, it appears, we will have Part III merger cases here at the FTC. So, I would say that there's a different role for merger simulation in each of those three channels.

So, the first channel I mentioned is the merger PI. I have some recent painful experience with that channel in the Whole Foods case. It -- I think it's real important to remember your audience here. We're dealing with a District Court judge. He or she is a very busy person. They are not an antitrust specialist. They're also going to be subject to severe timing constraints. They have a docket outside of the case that they're working on with you. A lot of it's going to be criminal. They have deadlines within that docket. And, you know, as we learned last summer, they are people who just may want to go to the beach the third week in August, and that's going to dictate their timing.

Additionally, they are going to be external constraints timingwise. One of the first things that the merging parties are going to say to them when they first meet them is, you know, our financing is going to fall apart in three weeks unless you, Judge, make a quick decision on this case. So, and that's a pretty uniform occurrence.

The other thing, final thing, to remember about the District Court judges is that they -- they're going to be risk-averse, and they're going to slavishly follow Supreme Court precedent, and the precedent that they are currently bound by, among them is Brown Shoe, and I'll just, you know, I'll just quote what they have to follow there, which is, you know, the proper definition of the market is a necessary predicate to an examination of the competition that may be affected by the horizontal aspects of the merger.

So, if we're looking at merger simulation as an alternative to structural analysis and the market definition analysis, then that's going to be an issue for pretty much all District Court judges, and it will cut against that analysis. So, that's the merger PI. Sorry to be so oblique, but I would also support what I'm saying with a quote from Judge Hogan, who's here in the District of Columbia, and some of you will be familiar with his name. He was the Judge who decided two PI cases in favor of the FTC. The first was the Swedish Match case, and the second was Staples in 1997, and he's quoted in this wonderful book, which I highly recommend to you. It's World War 3.0. It's actually about the Microsoft case. It's done by a guy named Ken Auletta, who writes in The New York Times, I believe. I'm sure it's a name you're familiar with. So, when interviewed about the Staples case, here's what Judge

Hogan had to say about his decision-making process in that case.

Judge Hogan recalls, "We had a lot of economic evidence in Staples. We had a lot of documentary evidence, although in that case, the econometric evidence that the Government had was not at all convincing to me." Sorry, folks. "I think the internal company documents were more convincing. That's why I stopped the merger." And then he went on to add, "a case with a judge or jury is won or lost on a handful of a few key points. You want to identify them early, marshal your evidence, protect them, attack the key positions of your opponent, and not get bogged down in a lot of detail, because fundamentally, at the end of the day, this whole case is going to get boiled down to a 35-page brief. At the appellate level, it's all going to come down to that." And we've seen that happen with the Whole Foods case.

So, the second channel I mentioned is the merger review process here at the FTC, where, you know, in a lot of cases, a closed decision is made or a divestiture is accepted, and as Mike's already explained, there are quite a few cases in the past ten years where those decisions have relied in good part on merger simulation done, particularly where the products involved were consumer products, whether there was reliable Nielsen data, IRI scanner data, and also I would add to that, from a lawyer's perspective, a lot of those were cases where there was a timing agreement with the parties where people -- where the economists had time to run models and do good work, and so they were informative in those cases, the merger simulations.

And then the final channel, I just wanted to talk about, was, as I see it, this emerging Part III merger litigation process here at the FTC, and that's something that is now increasingly coming to the fore, and the Part III rules here regarding merger litigation have been changed or are in the process of being changed. That's a process that's going to come through -- follow through pretty quick. There was a Federal Register notice published October 7th setting out the rule changes. And the new timing for merger Part III cases here at the FTC is five months.

So, as I see it, I'm not sort of acting on any special knowledge, I think that could be a real useful forum for merger simulation. The reason why I say that is because, you know, the FTC is the expert body, and we're charged with developing the law and policy relating to antitrust, and that might be -- you know, merger simulation or acceptance of merger simulation by judges here at the FTC and the Commissioners could be something that will increasingly come to the fore, although with what Mike said, I'm not so sure that the Bureau of Economics will be putting in as much.

And I would -- again, I would rely on another expert in making that observation. At least one of the Commissioners, Commissioner Rosch, as recently as June 2008, has alluded to the possibility in a speech that he made in which he -- he told his audience that he didn't necessarily see market definition structural analysis as a threshold issue in merger litigation.

He -- in fact, he went on to say, "I would suggest this is a mistake. A focus on market definition risks obscuring the ultimate section under Section 7 of the Clayton Act, which is whether a transaction is likely to substantially lessen competition, the answer to that ultimate question may turn on market definition, but it doesn't have to in all cases." So, I would see perhaps some hope there for merger simulation.

That's it for me.

MR. BAJARI: So, why don't we take one to two minutes each, if you would just like to follow up on each other's comments, and then we'll go ask the audience for some input.

MR. VITA: I'll just react to a couple of Aviv's comments. I think he raises a very fair point. You know, we're talking about merger simulation. I mean, sort of everything we do is a merger simulation. You know, I've talked about it in a very narrow sense in the way that we frequently use it here at the FTC, but I think you raise a good point, a fair point.

You know, when you do merger simulation, quantitative merger simulation, like we have been doing it, you know, it produces a point estimate of a price increase that allows you to go out and say, how well did that point estimate actually reflect what happened? Nobody ever writes a paper like Craig Peters' paper or Matt Weinberg or Dan Hosken's paper saying, well, how well does the traditional way of plugging market shares into Hirfendahls, how well did that predict the price increase? Because, frankly, it doesn't predict the price increase. It says, is there going to be one or not or, you know, maybe it might say is there going to be a big one or not, but it doesn't say, like, it's going to be 8.2 percent. So, it is a bit of an unfair comparison.

But I think that, you know, the bigger, harder question for people like me, for people like Gail, and for our Commissioners and for judges, you know, to the extent they entertain this evidence is not so much is it -- you know, should we do it or, you know, could it be valuable? It's what weight do you attach to it? And that's what I struggle with.

You know, I think it's -- I still believe that I think it's a worthwhile thing to do if you have the resources with which to do it, but in any given case, I don't know, well, how well does the Bertrand assumption fit? I don't know the answer to questions like that in the preliminary injunction world, in the Hart-Scott-Rodino world, where we're doing these things under a very compressed time schedule.

We have got to come up with an answer, and we have got to advise, you know, our decision-makers, you know, should you pull the trigger or not and how much weight did you attach to it? And I don't know the answer to that. And even if it's true, like you say, I mean, you are never going to get it exactly right in any given instance. You know, if you get it on average right, that's pretty good. I don't know that we're getting it on average right, but I think variance matters, again, you know, in terms of telling a decision-maker, how much weight should you attach to it? And as I sit here today, I'm not sure what answer to give.

MR. NEVO: So, I'll be brief, because I

basically, I think, took half of my previous five minutes to respond to the five minutes before me. So, just, you know, how much weight? I mean, I don't know. I mean, I think it depends on the case. I mean, I can tell you, you know, I think there is a difference when you do an academic work, you probably put a weight on one merger simulation of some form, but when you actually go and you look at a real case, I mean, I can tell you the cases that, you know, I've looked at.

I mean, the very first thing that you start is you do look at the data and do look at sort of -- I mean, you want to compute concentration ratios, but, you know, I've never found those sort of very enlightening, but try to sort of see maybe more reduced form sort of impact of things that you could look at. So, I don't know that there's a general formula, but -- and I think in that sense, the retrospective studies might be useful or -- no, not might be, will be very useful in helping us sort of get a feel as to what's a good assumption where, right?

So, a simulation worked well for a merger in ice cream that happened, you know, five years ago, and now you have a merger in chocolate, maybe they're close enough that there might be some similarities. Maybe not. I don't know.

MS. SLATER: So, yeah, when I hear the word "weight," I think of it as an evidentiary matter, and so, you know -- I'm sorry.

I just said when I hear mention of the word "weight," I'm thinking of it as an evidentiary matter, and the other wrinkle on looking at one merger simulation is that in the PI context or litigation context, generally, you are going to have another economist with a whole other way of looking at things, and also, almost by definition, they are going to come to the opposite conclusion. And so as a judge, you know, how do you deal with that and.

And one thought that I've heard bandied about is that, you know, it could be a timely thing for our ALJs here at the Commission to have their own -- to have the power to hire their own economic expert who can -- who can do just that, who can weigh both expert opinions and look at each and attach weight on behalf of the judge and advise the judge on that issue, so...

MR. BAJARI: So, let's go to our audience next and see if we have any questions.

Carl?

AUDIENCE MEMBER: So, the question is, how much merger simulation has really affected decisions in the Commission, the investigative process and decision to challenge or really seek a remedy, I guess? My sense is in the courts, it has had very little effect, in large part or -- because it's so opaque, and the experts come in, and it's quite sensitive, right? The functional form that you use and also to other things, judges, very hard for them to sort out. So, has it really mattered or not?

MR. VITA: I'll take it. I'll try to answer that. Well, let's start with, you know, how does it affect the decision-making within the Agency? I mean, really, I'm not a decision-maker, so I'm not really the person that can answer that. You know, we hope that it influences people, but, you know, again, what weight do our Commissioners place on it when the Bureau of Economics forwards a recommendation memo that includes, you know, a simulation exercise as part of the evidence that we think is relevant?

I think some Commissioners -- you know, there's variance, but I really don't know. I mean, I -- that's -- it's not a question I can answer. I think there -- with certain people, it goes right over their head, and they attach a weight of zero. They're like Judge Hogan. I mean, Judge Hogan's comments that Gail read are disquieting, because they don't really get to the issue. You know, again, we're talking about merger simulation in a very narrow sense of estimating functions and, you know, doing some oligopoly simulation.

The stuff that was done in Staples, that's about as simple a quantitative piece of analysis as an economist can construct for antitrust, and if that's too hard for people, we've got to find another line of work, because, you know, I don't know what our contribution is, at least, you know, in doing quantitative analysis. You know, when you get to something harder like, you know, the kind of stuff Jerry Hausman does where, you know, explicit assumption about a particular oligopoly model and that sort of thing, my guess is -- I don't know how often that's actually been presented in court. I don't know that we, at the FTC, have ever presented such an analysis in court. The Department of Justice may have. My guess is, you know, my guess is the typical judge is not unlike judge Hogan. That's my guess.

MR. BAJARI: Any other comments?

MS. SLATER: Well, I think some judges or at least one that I'm aware of has accepted critical loss as a simple story and relied on it quite heavily.

MR. VITA: Let's not go there.

MS. SLATER: But this panel is not about

critical loss, so...

MR. NEVO: I just have actually sort of a related -- I mean, it's almost a question. I mean, we're envisioning sort of -- you know, here, there was sort of a -- I guess a quota sort of saying there was no effect, but what would have happened if one side comes up with an analysis and the -- I mean, so we should not be mistaken by sort of having kind of the equilibrium phenomena of both sides coming out and cancelling each other, versus if one side came with a very detailed model --

MR. VITA: There is actually a data point on that. It's Whole Foods. That's exactly what happened in Whole Foods. The FTC, through its expert, Kevin Murphy, presented a Staples-like analysis of the likely effects of the transaction, you know, looking at how entry and exit events affected prices in geographic markets, and the witness for the other side didn't do anything like that. I mean, you know, did different stuff, but he didn't do that. And, you know, we know what happened there, so --

MS. SLATER: I think that witness even went on to describe Kevin Murphy's work as some of the most sophisticated modeling he had ever seen in his entire career.

MR. VITA: He said it was really hard.

MS. SLATER: There you go.

MR. VITA: Probably shouldn't talk about him, but -- yeah. So, I mean, no, you would think, yeah, something trumps nothing, but not always.

AUDIENCE MEMBER: I actually had a question that related sort of to the Whole Foods. I've worked some with sort of the very narrow merger simulation models. Something that I've pulled out of it that I didn't really appreciate before I got into them was that what really seems to matter is the cross-elasticity or diversion ratio between the merging products, and Carl's done work on that, Dan O'Brien, and Abe Wickelgren have a paper, too, that really highlight the diversion ratio as an important piece of evidence.

And that came up in Whole Foods, and there were discussions of what seemed to me to be very large diversion ratios that, I think, the judge characterized as not being that large. And I'm just wondering what sort of evidence we, as economists or attorneys, could bring in, if not with sort of structural oligopoly models? Is there empirical evidence that we could develop that would sort of help us to highlight to judges, you know, here are sort of the diversion ratios in particular markets that -- that would be problematic.

Is the best way to go at that a structural model? Is there other empirical evidence we could develop that would help put those numbers in context? Does anybody have any opinion?

MR. VITA: Well, I mean, you know, we don't want to get into a session where we complain about what the judge did or didn't get right in that case, but there, I mean -- I mean, I've -- my view of that case is we had a pretty simple, straightforward story, and it was one where the -- you know, if you view sort of merger simulation more broadly defined, there we did it, again, with some sort of reduced form, Staples, that exercise that Murphy carried out. It was, I thought, a great complement to an abundance of traditional kinds of antitrust evidence that we got from documents and testimony and that sort of thing.

Where I -- where I think, you know, when you read, you know, the decision in that case, I think -- I think to me, it betrayed a fundamental lack of understanding of sort of the basics, just the -- you know, why is the diversion ratio or the cross-elasticity of demand important, you know, in trying to forecast or trying to predict what the competitive effect of a given transaction is likely to look like? And that -- you know, the solution to that is to, you know, have these guys go take Economics 101, I mean, which judges sometimes do. I mean, there are, like, law and economics programs to try to instill in them sort of the basics.

But, you know, so I mean it is -- it is a little -- you know, it is a little depressing, because, I mean, how are you -- you know, the idea that somehow an elaborate, sophisticated merger simulation exercise based upon sort of state-of-the-art techniques, the kind of stuff that Aviv and Pat do, you know, what role is that going to play? Well, when a judge doesn't even understand the fact that, you know, a high diversion ratio between the merging parties, other things equal, means there's a pretty high likelihood that prices are going to go up. If you can't grasp that, you know, I don't -- you know, I don't know what to do. Do more research or something, you know, but it's -- it's -- you know, that's very depressing.

What do you think, Gail?

MS. SLATER: I can't add to that, Mike. Sorry.

MR. BAJARI: Aviv, do you have anything to say?

MR. NEVO: Not -- maybe not directly sort of on the real policy thing, but, you know, I think ultimately -- and, you know, that's kind of pushing a little bit, you know, going back to the retrospective studies, I mean, I think if we look -- take a broader view of these studies, rather than to, you know, beat on any particular method, right, ultimately what we can get from those is some sort of a database that will give us, you know, if you want some sort of a mapping between diversion ratios and actual outcomes, right?

So, I mean, I think if you have a translation to show to a judge, show him, look, in the past ten years, here are sort of mergers we thought these were the diversion ratios, and these were the outcomes. Now, whether they match exactly some Nash Bertrand prediction or not, they're not going to match exactly, but if you kind of show that there is sort of a systematic relationship there, you know, I think at that point, it does become relevant, you know, the diversion ratio.

Now, if there isn't a systematic relationship, then maybe diversion ratio isn't important. I mean, then we have to figure out why, right? But if it really is relevant, we should be able to see it in the past data. And, you know, again, not focus on the exact specifics on did we get Nash Bertrand or not Nash Bertrand, but is there sort of a general mapping between, you know, the estimated diversion ratios and what happened? And I think that's the kind of evidence that you want to sort of put forward, and once you can establish that, I think it would be easier to make the claim.

MR. BAJARI: Do we have any last questions from our audience here? Yeah.

AUDIENCE MEMBER: So, obviously, other -- I mean, two other kinds of evidence that are sometimes used, as has already been mentioned, the company documents from the merging parties, but also, particularly in wholesale mergers, is some surveys of customers, so, say, retailers or -- now, there's obviously issues of truthful revelation that go on when you handle those documents, but is there a sense about how accurate those documents tend to be about what happens ex post?

So, you know, if you do a survey of customers, you know, are customers tending to get it right? Because, you know, that's somewhat -- that's going to be somewhat informative about, you know, is there some thing which we as outsiders, when we at these markets are missing that people inside the market always understood, or is it the case that actually there's just a lot of things change over time that are simply going to be very hard for anyone to predict?

MR. VITA: I would say -- respond quickly, I mean, I think it's a very good question. I don't really know the answer to that. I mean, we don't know -- you know, I know the European -- a lot of the European antitrust agencies, formal surveys are actually a pretty important part of the decision-making process. That's something we haven't done here, and it might be something we might want to think about. I don't know that -- it would be interesting to note if the people who do that kind of analysis and gather that kind of evidence have ever done any kind of ex post evaluation to figure out, how well does this work? So, that's a good question, but I -- you know, I really can't say any more about it, but I think it's an interesting idea.

MS. SLATER: So, if I understood your question correctly, you said there were two sources, I think, identified as the polling, and then there's the internal company documents, and there was -- there was a poll done in the Whole Foods case, not by us, by Whole Foods, and that was -- that was successfully, effectively Dauberted by the FTC, because there were flaws with the survey instrument and with its execution, and so, you know, there are issues here in this jurisdiction with surveys and Daubert that don't exist in Europe, because it's an administrative process over there, and so that they're more problematic here.

With regard to the internal documents, to me --

and this is public domain, so I don't think I'm giving away anything here -- to me, some of the most compelling internal documents that I reviewed and relied on in depositions were the Wild Oats documents where you had actual diversions, the past being prologue of where a Whole Foods had opened close to a Wild Oats store and what happened to sales at those stores, and you had actual diversions of, you know, sales losses from the Wild Oats store, diversions -- I wouldn't get too technical -- of up to 30 percent of revenues, actually more than 30 percent of revenues in some cases, so...

MR. NEVO: Again, sort of speaking from a different angle, the thing that always sort of struck me is seeing -- you know, sometimes, you know, talking not in the context of cases, but how they would actually put weight -- you know, companies would actually put a lot of weight on these sort of surveys that, you know, if you were given that data, you would look at it and say, "What can I do with this? I mean, this is pure garbage?"

Yet they're spending millions of dollars, and you can ask them, why don't you go -- you know, there is this company, Nielsen. You must have heard of them. They collect all this, you know, actual transaction data, or they have this home scan data, you know, why don't you buy that? I mean, it's not going to cost you anymore.

Yes, you might have to pay for, you know, for an economist or a statistician to do some analysis, but sometimes, you know, they won't go to that. And that's sort of -- I don't know if you want -- that's the miles an hour error image of that, not even the -- you know, the policy aspect. I don't even know inside, just from a profit maximization, why do they actually think that one form of data is useful if they're not willing to use the type of data that we think is useful?

MR. BAJARI: So, would our panelists like to take one minute and give some final thoughts each, and then we'll go have some coffee?

MR. VITA: Okay. Well, I'll just say, you know, I continue to think that the information generated by these kinds of quantitative analyses are useful and important, but I really do encourage people who -- you know, who are interested in this kind of stuff to continue to do the kinds of retrospective papers to get us a better idea of how -- you know, how we can refine those tools and make them -- make them work more efficiently.

MR. NEVO: I pretty much sort of agree with everything, and I think, you know, what we need to do is, again, to continue and bring out good economics to try and address these types of very hard questions.

MS. SLATER: And I will add, as the lawyer in the room, you know, you guys are really smart guys. What you do is very valuable. There is a place for it. Know your audience and cut the cloth to fit the audience.

MR. BAJARI: All right. Well, I'd like to thank all three of our panelists. This was a lot of fun.

(Panel Session One concluded.)

**PAPER SESSION ONE: ESTIMATING DEMAND AND PRICING STRATEGIES SESSION**

MR. STERN: Okay. So, we are delighted to start the first paper session of this conference, and this actually -- and our session is on what says estimating demand and pricing strategies, and I think some of the papers are actually on those topics. So, the -- we're going to start with Alan Sorensen of Stanford on the welfare effects of ticket resale.

MR. SORENSON: Thank you.

Okay, thanks. It's good to be here. This is joint work with Phillip Leslie, and it's funny that Scott talked about repugnance. I didn't anticipate that, and the first bullet point on my first slide is about repugnance.

We're talking about ticket resale here, and I think it's good by pointing out, as many of you know, that ticket resale is something that's looked on with some measure of disdain, pretty much anywhere you go in the world, to varying degrees in different cultures. And, in fact, in the United States, I think we're the most generous of the countries that I've looked at toward ticket resellers, but in general, this is a controversial activity, but it's been growing substantially in the United States due largely to the onset of online markets for ticket resale.

The other interesting trend that there is, at least in the United States, is a trend of deregulation. There are five states that last year basically repealed their anti-scalping laws. So, the legislative trend is towards making resale okay, whereas before, there were many restrictions on it.

So, I said that this is a controversial activity, but, of course, most economists think of resale markets as good things, and it's tempting to think of resale markets as being unambiguously good things, because the transactions that are taking place are voluntary transactions, they are reallocating goods to high-value customers or consumers, which is exactly the kind of thing we want markets to do, and these two things suggest that these are Prado-improving (phonetic) trades.

The problem is or the complication is that that presumes or it takes as given the initial allocation, and the thing is, the initial allocation in a resale market is endogenous in the following sense in that who buys or the decisions that get made in the primary market are endogenous to what's going to happen in the secondary market in the sense that they depend on what people expect is going to happen in the secondary market. So, that complicates things.

And in particular, the existence of a resale market attracts brokers, okay? So, they're going to create some value, perhaps, by providing liquidity in the market, but they might also capture some value. So, we argue that the welfare consequences of resale markets are actually a little bit more complicated than you might think at first blush.

The first -- and I think the first order thing that we want to talk about in that paper is the reallocative effects of resale. These are the good things, right? This is what increases the size of the pie, because it's reallocating tickets to people with the highest values. But there are other complexities. For example, allowing resale is likely to increase the prices paid by the people who actually attend the event. It generates some profits for brokers. It can generate profits for consumers, as well. Professional brokers aren't the only ones arbitraging in this market, right? Individuals like you and me are also playing around in this, and there's evidence of this in our data.

Resale can cause people to buy in the primary market who otherwise would not have, that is, anticipating the potential for profit in the resale market, some people might buy when they otherwise wouldn't have, and it could also have the opposite effect on other people. It might cause people to wait, thinking, why buy now when I know there's an active resale market and I can just wait and buy a ticket closer to the event?

It also increases competition for tickets in the primary market, and it might -- it might make it so that more people -- there is more intense competition early on, right when the tickets go on sale for an event, you might get this sort of mass onset of buyers trying to get tickets all at the same time, and in particular, it might be costly to get yourself early in the queue if people are trying to get those underpriced tickets, and the costs incurred should probably be weighed against the benefits of reallocation, okay? So, this is another complexity.

And then, finally, if you endogenize the primary market prices, the presence of a resale market might change the way the prices get set in the first place, okay?

So, the first of these pieces is the pie increasing piece, and the rest are all about how the pie gets shared, and we want to talk about all of them in this paper. We want to, but we won't. In particular -- I mean, this is a lot to talk about. We're going to try pretty hard to see what we can say about the first point, because we think that's sort of the first order interesting thing. The others we're going to have something to say about, how the pie gets shared among the various market participants, but numbers 7 and 8, in particular, we're not going to have much to say about yet, which is unfortunate, because 7 in particular, I think, is a very interesting question, and we have some ideas about what we can do to address that, but at least in the current draft of the paper, there's pretty much nothing we can say about that particular point.

So, what we are trying to do, then, in this paper is, you know, talk about how much resale increases aggregate welfare by this -- by reallocating tickets to high-value customers. We want to say something about how that pie is shared, who's winning, and who's losing. We also want to be able to say things about what would happen if we could change fundamental characteristics of resale markets.

So, for example, if we could exogenously lower the costs of transactioning in resale markets, what's that going to do to the payoff to sellers, consumers, and brokers? And the reason this is an interesting question is you can view the Internet as having done exactly this, right, just lowering the costs of transacting in resale markets.

And then, finally, we're going to be able to say some things about how resale markets would look different if primary market pricing patterns were to change, and we're going to do this by estimating a structural model using very detailed data on 103 rock concerts from the summer of 2004.

I'll tell you about the data in a minute, but let me just give you a preview of what the model is going to look like. It's going to be a two-period model. The buyers in the market will either be brokers or consumers, and the distinct -- the technical distinction between them will simply be that a broker in our structural model is a buyer who has no utility from actually attending the event, okay? So, they only buy with the anticipation of reselling.

And a key part of the model and what makes it interesting and also what makes it complicated is that it's a rational expectations model. So, when buyers make their first period decisions, they have expectations about how things are going to play out in the second period, and those periods -- those expectations are going to be correct in the sense that the decisions that they lead to deliver equilibrium outcomes in the second period that turn out to be, on average, correct or, on average, consistent with the expectations. So, that piece of the puzzle turns out to be complicated, but it's important, because in the end, that's what ties the primary and secondary markets together.

In our model, there are two sources of friction in the resale market. There are transaction costs, which we allow to be different from brokers versus consumers, and there's also friction in the sense that we clear the resale market through a sequence of auctions, and we sort of exogenously impose that there's random participation in these auctions.

Now, this is a data-driven assumption. We wanted to fit a certain aspect of the data, but it is an assumption that sort of not every -- the market isn't cleared in an Al Roth sort of way, where you just sort of line up the people with the highest valuations in order of valuation, and you line up the people with the lowest reservation values, and you just clear the market. It's a much noisier, more friction -- frictioned process than that, but there's a reason for it in the data, which I'll try to mention when I get to it.

And also, there's uncertainty about the overall level of demand in the primary market; that is, the buyers in this market aren't perfect predictors of the demand for a given concert, and the reason we wanted to do that is that there's plenty of evidence in the data that some of the speculators actually get burned. Even the professional brokers often buy tickets that they end up having to unload at below face value. So, it must be the case that there's some uncertainty, ex ante, about the strength of demand for a given event.

Okay, so like many empirical projects, this is one in which a lot of effort went into obtaining the data and then working on data to clean it up and get it ready to analyze. I'm going to spare you the details of that process and instead just show you a picture of the data. We have data from Ticketmaster, which is the primary market seller for these concerts, and we simultaneously have data from StubHub and eBay, which are the two leading online ticket resale sites. So, we're seeing what gets sold in the primary market, and we can see, in parallel, what's getting resold in the secondary market. So, this is the picture of our data for one of the 103 concerts in our data set, Kenny Chesney at the Tacoma Dome in Washington, on June 17th in 2004.

So, if you look at the horizontal -- everyone's looking at a different screen, so I can't point at a screen, but maybe I can -- yeah, okay. So, if you look at these -- whoops. Let me back up.

If you look at the horizontal lines here, those are actually a bunch of points. Each point is a seat in the venue, and those are the prices -- what's plotted on the vertical axis is the price at which those tickets were sold in the primary market. So, you can see that this concert had 20,752 seats in the venue. They're ranked from worst seat quality to best. There were three pricing tiers for this event, basically going from just under $50 to around $70. So, basically, this event sold out. That's why these horizontal lines sort of go all the way across.

And then all of these other points are tickets that were resold, either on StubHub or on e-Bay. The squares are broker resales, tickets that were resold by brokers. And the circles are tickets that were resold by sellers who appear not to be brokers, and they appear not to be brokers because they only show up in the data a couple of times, okay? We call a broker somebody who shows up in our data as having resold ten or more tickets during this summer, okay?

So, a couple of things you can see just by looking at the data. Most of the resale activity is concentrated at the upper end of the quality spectrum, which is not surprising. There's a big premium for the very, very best seats. So, this is a concert where the tickets were, in general, underpriced. You can just see that from looking at this picture, but this is especially true of the very best seats. The front row was vastly underpriced, and that's typical of concerts in our data set.

But you can also see that here and there, a ticket gets sold below face value. It doesn't happen as often for this concert, but across our entire data set, you'll see that it happens more than you might have guessed. And another interesting fact here is that there's some clustering of resales at the upper end of the quality spectrum within a pricing tier, okay? That is, if you look at the sort of second pricing tier, the best seats within that tier were underpriced, and so there's -- you can see that brokers were gobbling those up and reselling them. So, that's a picture of our data, and I hope it gives you a good sense of the level of detail we're working with when we move on to estimated the structural model.

The quick summary statistics, I'm going to speed through -- there is just a couple things. The average markup is just over $20; in percentage terms, the average markup is 40 percent. But it does happen -- let me scoot down here -- about a quarter of tickets -- a quarter to a third of tickets are actually resold below face value. So, that's either speculators that are getting burned or people who are unloading tickets that they ended up not being able to use and having to unload them below face value, but that's sort of an important fact that I don't think we anticipated going into this analysis.

The timing of transactions -- in the primary market, almost all the sales occur right after the tickets go on sale. I'm not going to explain this picture, but we used this picture, if you look at that paper, as a way of justifying looking at a two-period model instead of a multiday model. The fact of the matter is that there's some overlap between these markets, and there are some interesting dynamics, as you'll see in Andrew's paper, which comes up next, but we're pretending that those dynamics are uninteresting so we can have a simpler model to estimate.

Okay. So -- okay. Just quickly, an overview of the model. So, in the primary market stage, we've got brokers and consumers arriving in a random sequence, and they all have expectations about the resale value of each seat, and they make their decisions based on those expectations, and there's heterogeneity in consumers' willingness to pay for a ticket. Of course, that heterogeneity, that makes it so that there are gains from reallocation, right? They're arriving randomly, and the low-value guys might get lucky and arrive early and get a high-value seat, and that's what gives rise to gains from reallocation in the resale market.

And in period two, the way we clear the resale market is we take the allocation that obtained in the first period, and then we have a sequence of auctions, starting with the highest-quality seat, we have the holder of that ticket holding a hypothetical auction and then randomly we select from the pool of potential buyers, and we conduct a second price auction, okay? So, that random participation is a major source of friction in our model, and we do that sequentially, starting with the highest quality ticket and proceeding on.

We assume that there's no option to return to the primary market. So, we assume away the idea that you might buy a ticket in the primary market, sell it on the secondary market, and then go back and buy another ticket in the resale market. So, there's no buy-sell-buy behavior. There's some evidence that that occurs occasionally in the data, but we're assuming it away for simplicity. I point it out because I don't think it's an entirely innocuous assumption. I think some of our conclusions might change a little bit if we could fix that, which we're trying to do.

Okay, I'm going to skip that slide.

Yeah, so, I'm going to -- that's okay, because I'm going to go quickly through this. So, the -- the preliminary estimates are boring. So, I'm not going to talk about those, but one thing that's a little interesting is that we're estimating very high transaction costs for consumers, transaction costs on the order of $70, and that's the cost of going and selling a ticket on eBay. Now, you might think that that represents a number of things, and I know Mary and Alan Kreuger have talked about there being endowment effects, and in our analysis, an endowment effect, if it exists, is just embodied in this transaction cost. Nevertheless, it's worth thinking about is that transaction cost implausibly high? We think it's not implausibly high, but it's something to highlight.

So, now I'm just going to give you a -- there's no way I can go through all these numbers, obviously, but I just want to give you an idea of the kinds of counterfactual analyses we're wanting to conduct once we have estimates of our structural model. We can sort of simulate what was going to happen in the base case, under the current regime, and then compare it to what happened, say, if we zeroed out transaction costs, right? So, if we take the parameters for transaction costs for consumers and brokers and we just set those equal to zero, we can then simulate what these resale markets would look like and what the outcomes would be for all the various players.

So, for example, if we do that particular comparison, the number of tickets sold in the primary market is largely unaffected, but the fraction of tickets that get resold in the secondary market goes up dramatically, not surprisingly, and then if you sort of go down here, if you look at the bottom line in that red square, total surplus goes up substantially if we zero out transaction costs. So, eliminating a friction in the resale market increases the size of the pie. So, that's sort of what you would expect to see.

Some other kinds of comparisons we can do, we can compare zero transaction costs to a market in which there's no resale at all. If we could shut down the resale market altogether, how is that going to affect total surplus and the share of surplus? There's another thing that we can do, which is that -- something that I haven't talked about at all is that the reason these resale markets exist is that the primary market prices are so badly constructed, right? I mean, the pricing in the primary market is so coarse that that's why all these tickets are getting resold.

So, a question that's natural to ask is, what would things look like if the primary market sellers just did a better job of pricing their tickets in the first place? So, that's sort of this comparing column 1 to column 8 in this counterfactuals table, what we're trying to do is say, suppose we took the top 10 percent of each venue and we priced it optimally in some sense? So, we made the primary market seller a bit for sophisticated, what would that do? And, you know, what we're showing here is that it would increase the producer's revenues, you know, not dramatically, but it would definitely increase the seller's revenues, but the idea here is that we're trying to quantify the gains from more sophisticated pricing.

And moreover, another interesting feature of this and something that I've talked to people at Ticketmaster about, because, of course, they're now pushing to have more auctions in the primary market stage, to the extent they do a better job of pricing markets -- pricing tickets in the primary market, we should expect to see resale activity diminish, right? We don't have -- our model allows us to simulate that, but given that Ticketmaster is now conducting a lot more auctions, in principle, there's data out there to test that directly, whether holding an auction leads to less resale in the secondary market.

So, you know, this paper is still a work in progress, but the conclusions we're drawing at this point are that coarse pricing in the primary market is really what's driving resale activity.  It's not just that tickets are underpriced in a general sense.  It's that these guys are only setting two or three or maybe four prices for tickets in a 10,000-seat venue, right? So, that's a very coarse pricing structure, and the resale market is largely about undoing that.

And then, finally, resale markets are redistributing surplus in some fairly subtle ways, but the numbers we're coming up with at this point suggest that the observed levels of resale activity aren't large enough to make a big difference in overall welfare; that is, marginal changes to the amount of resale activity aren't going to have a big impact on total welfare, but large reductions in transaction costs, like the reductions that have been affected by the Internet, I would argue, could generate substantial gains in total welfare, could increase the pie substantially.

And then finally, an interesting point that I think is sort of obvious when you think about it, but you have to think about it first, is with near frictionless resale markets, consumers who actually attend the events are probably the biggest losers. I don't mean they're a loser for attending the event. What I mean is that they actually, in a frictionless resale market, if you end up going to the event, it's probably because you ended up paying a very high price for your ticket. So, if you -- a lot of you will want to think about what's the problem for the primary market seller and why are they underpricing their tickets. In a market with frictionless resale, artists are going to playing concerts to people who aren't sitting on a lot of surplus when they are sitting in their seat, right? They basically -- most of their surplus got extracted in the process of them obtaining the ticket. So, I'll stop there. Thanks.

MR. STERN: Our discussant is Mary Connelly Pray.

MS. PRAY: Okay, hi.

So, I'm here to discuss Alan's paper that I really enjoyed reading, because it -- well, it's a good paper, I think, though not completely finished, but on the way to become a really good paper, and also because it relates to work that I've been doing jointly with

Alan Kreuger, and so just to highlight, you know, what was in Alan's paper, but I'm not -- you know, there is no need for me to redo his presentation, he looks at the social welfare impacts of ticket resale, and more specifically, who wins and who loses when there is ticket resale, and I think his structural model estimation allows him to then do different -- you know, run different specification and look at what would happen under different scenarios, both from the point of view of the producer, the brokers, and the consumers. And I think that's really interesting.

Now, the model has, you know, two types of agents, the brokers and the consumer, and I think it's also really interesting that he puts in there that the -- you know, you can buy on the primary market and then attend or resell, or you can also buy on the secondary market, or you can just not buy at all, and I think the interesting part is also that a regular consumer could buy a ticket and then resell it, and I think that adds well to the model.

Now, the data -- the data that Alan and Phillip used are coming from Ticketmaster, eBay and StubHub, and though he didn't mention it very much in his presentation, you might be concerned that the resale part of the market that he has only come from -- the data only come from eBay and Stubhub, which are clearly not the only two players in the market.

Now, forget in 2004, but, you know, now a days when we think about it, people sell their tickets on Craig's List, brokers have their individual Web sites. There are other Web sites, other than StubHub, like TicketsNow and things like that, and also, some people -- there is still the shady scalper at the door selling the ticket, and people might be selling tickets to their friends. So, there are all those transactions that are not part of the data that Alan and Phillip use in the market. And, you know, people might be concerned about that fact.

Now, I'm going to -- you know, the bulk of my input to this discussion is going to be to present some preliminary results that Alan Kreuger and I came out from -- that we're taking from a national representative survey of concerts that we ran from August to October 2006, where what we did -- and this is still preliminary work that, you know, we're hoping to finish soon -- what we did is we surveyed concert-goers inside the venue at different -- at concerts, at 30 different concerts, that were randomly selected across a universe of concerts during those three months, and we asked them, you know, what seat did they have; where they bought their ticket, you know, how did they get their ticket; how much they paid for it; you know, some reasons why they might have gone through the secondary market, if they did; you know, and other questions like that. And using this data, you know, we can now compare how the eBay and StubHub data that Alan uses with our data.

So, their overall resale rate is 4 percent, and our overall resale rate is 10 percent. Note that that is good, because, Alan, correct me if I'm wrong, but in your estimation, you're implying that -- I mean you have a parameter there that says that your eBay and StubHub account for 50 percent of the data. So, that's, you know, roughly correct. In our estimate, the market shares of eBay and StubHub are 31 percent of the total.

Now, also note that Alan's data is from 2004, and our data is from 2006. So, there might be also differences in there, which the market has evolved in the two years, and we know this market changes very fast.

The average ticket price in the primary market, in Alan's data, is $83; ours is 81. If you look at the price for the seat that they mentioned that, you know, they're sitting in, but the average price that they reported paying is $88. So, again, roughly the same ballpark.

The average resale price is -- from this paper is $111; ours is 122. Again, a similar ballpark. Average list price of the resold ticket is $89; ours is 91. Again, this -- this lines up pretty well. And we also find that if you compare the second item with the fourth, the average ticket resold is a better ticket, it's worth more than the -- you know, than other tickets, tickets that were not resold. So, the best seats are resold, we did find out in our data, and the average markup is 39 percent, whereas in our study, we found a 36 percent average markup. So, again, that seems really well -- to be matching up really well.

Now, a point that Alan mentioned is part of the results -- as part of the results, he finds that the consumer's transaction costs are $63 versus the broker's cost of $12, and it's true that this might seem like a high number, and a partial explanation that he gives, you know, circles back to us in the endowment effect, and that is a point that we also find in this survey, which is that the consumer's valuation of tickets increase after purchasing them.

And what we did to try to test if that effect is true is we asked people two questions -- well, we asked two different questions. We had two types of surveys, and so we randomly asked half of the people one question and the other half another question. And the questions were, would you have bought your ticket if it would have cost you $300? And the second question, if someone offered you $300 for your ticket, would you have sold it? And only 11 percent of the people said that they would have bought the ticket for $300. Now, since some of the answers on the paper were clearly, because I paid $500, and now I -- 47 percent of the people said that they would have sold their tickets for $300, and so we can see here an example of this endowment effect.

And so in conclusion, because I have to stop, I think it's a very rich model that combines various insights from the paper and uses unique data from both the primary and secondary markets, and as Alan mentioned, an interesting way forward would be to think about the pricing in the primary market and how that has changed with the evolution of the Internet markets.

Thank you.

MR. STERN: Okay. What I thought we would do -- and, Chris, tell me if this is a bad way of organizing -- is maybe do one or two questions after each paper? Is that fair? Where's Chris? Okay? Okay.

So, Michael.

MR. BAYE: A quick question for Alan. I was just wondering, in the overall welfare effects story that you're telling about secondary markets making consumers worse off, and I think the welfare effects were pretty small in the counterfactual experiments you did. Can you identify whether any of those welfare effects are stemming from something you might think of, like double marginalization?

Obviously, in the Kenny Chesney ticket example where you're selling out anyway, there is not any reduction in overall number of ticket sales, but in the overall sample, are the welfare effects being driven by the fact that the double markup leads to fewer tickets being sold ultimately, or are they selling out?

MR. SORENSON: That's an interesting question. I mean, certainly -- okay, first of all, the factual response to your question is that, no, not all of the events in our sample sell out. I think roughly two-thirds do. It's in the summary statistics table. Most of the concerts in our sample, by the way, are fairly large concerts by big artists. It's not a random sample.

So, for at least half, probably more like two-thirds of the data, like you said, this isn't a relevant issue, but I think for the ones where the -- where the primary market tickets don't sell out, it is kind of an issue, right, because you're going to end up with -- I think we hadn't thought about it in this way in the double marginalization terms, but certainly we're thinking hard about how the presence of a resale market leads to fewer sales in the primary market. So, yes, I guess we haven't really thought about it in those terms. We will have to take a look at that.

AUDIENCE MEMBER: Sort of a tangential question about going to concerts and these events is ultimately a social experience, and so you might wonder, you know, your neighbors are going to matter for your experience. So, have you thought about when you price these things, you sort of worry about who's buying the tickets. So, I went to a Redskins game Monday night against the Steelers, half the audience was Steeler fan, and the next thing they were saying how that was a bad thing to happen, right, because the Redskins weren't, you know, being fully cheered on.

So, for concerts, I've heard similar things where, like, one band just had one price for all tickets in a stadium, because they wanted the most enthusiastic people to show up. So, I mean, have you thought about this at all in your research and other papers?

MR. SORENSON: Yeah, well, like I said, we're sort of -- we're trying to steer clear of explaining what happens in the primary market and why they price in the way they do, just because we're already tackling a lot in this paper, but there are tons of really interesting questions about why the primary market sellers do what they do.

But along the lines of what you were saying, one interesting point, I think, is that, you know, there's this tend to deregulate. There's this trend to repeal anti-scalping laws, and most economists would say that's probably a good thing, right? These are markets that are -- you know, they're voluntary exchanges. It's probably increasing total surplus. So, why bar them?

But one way to think about the legal argument for anti-scalping laws is that you're trying to protect the right of the seller to choose who gets to come to the event, and many of these artists care who comes to the event. The Redskins want mostly Redskins fans at their event, and if there's a way, through their pricing, that they can achieve that that would be undone if you allowed resale, then maybe -- you know, maybe as a society, we decide that we want laws to protect that right.

Certainly as universities, we don't want there to be resale markets for admission, right? We want to be able to choose who gets to sit in our classroom, right? And we like to protect that right, and maybe we want to apply similar rights to sellers of event tickets. I don't know, but that's one way to think about the problem.

MR. STERN: Okay, and our next is going to be Matthew Sweeting, and let me get up here. Okay, I'll let you figure out -- he's going to be presenting -- well...

MR. SWEETING: Okay, thank you. So, I'm going to be talking about the research which looks at the dynamics that happen in secondary markets for Major League Baseball tickets. So, this is the part of the paper which Alan decided to extract for his book. So, basically I'm going to be doing two things: Firstly, describing the dynamics of prices that we see, and they're going to be very stark and striking; and then secondly, to be testing kind of explanations for why what we see happens is the equilibrium outcome.

Now, I'm going to be looking at Major League Baseball tickets as an example of perishable goods, right? So, a ticket to a game is perishable in the sense that after the day the game is played, the ticket is effectively worthless. Now, another characteristic which affects some of the analysis is there's also a characteristic of fixed-date consumption, right? And by that what I mean is independent of when you actually buy the ticket, you can only go and enjoy it on the day the game is played.

Okay, so the theoretical models which in the end I kind of use my work to kind of illustrate are partly driven by what we see in the revenue management literature for how people should price perishable goods. So, the basic theoretical structure in these models is as follows: So, think about there being a seller who has a fixed number of units to sell before a certain date. We're going to assume that there's no commitment, so the seller can continuously vary the price in response to the time left until the game and how many units they have left.

There's a very simple demand structure. So, consumers arrive randomly, and we're going to assume that they can't -- in the simplest models, we assume they can't delay purchasing. And we're going to assume the demand parameters are constant over time.

So, what we see in these models is the optimal price at any point which the seller wants to set is going to reflect the probability that a sale today causes the seller to forgo a sale in the future, because they won't have a ticket left when other sellers arrive, right? So, what this means is that the fewer units you have left, the more likely a stock out is going to be in the future, so you want to set a higher price.

On the other hand, when there's less time remaining, there's less opportunities for other sellers to show up, and, therefore, that's going to tend to reduce the price at which you want to sell, okay?

So, a robust -- what's been described as a robust prediction in this model is the expected price -- the price you expect to observe should be falling over time. So, falling is the moment when the goods are going to perish approaches.

So, there's been actually -- I mean, obviously, revenue management models are widely used by firms to decide how to price products, but there's been relatively little work actually trying to test the motivations identified in the literature for how you should price are actually being used in practice. And when people have looked at this, for example, in airlines, you tend, for example, to reject the declining price prediction.

In airlines, at least as a couple of pretty obvious explanations for why that might be the case, so you might think that towards the end, consumers with more inelastic demands tend to turn up, say such as business people, and that's going to tend to provide the airline with an incentive to increase prices close to the date of departure, on the other hand, there may also be commitment incentive, which the simple models abstract away from.

So, a commitment model, you may -- a firm would be thinking, well, if I tend to cut prices over time, that's going to cause consumers to wait in the future, and maybe to prevent future waiting, because I'm going to win throughout with these consumers repeatedly, I want to, say, commit to having a flat price schedule and maybe an increasing price schedule.

Okay, so I'm going to be looking at secondary ticket markets for Major League Baseball tickets, right? Now, these are going to be, I think, a nice example to look for kind of the revenue motivation, the declining price, for a couple of reasons.

So, firstly, sellers in these markets are very small, and this is actually one aspect where my data is going to differ a bit from Alan's data. So, in Alan's data, it's for concerts. A lot of secondary market sellers are actually fairly large brokers. In my data, a lot of the secondary market sellers are going to be very small, and what they are is a season ticket holder. If you own a season ticket, you have the right to go to 81 games. Even really loyal supporters don't want to actually go to all 81 games. So, they sell their couple of tickets that they have for the games that they don't.

So, when you look at HHIs, for example, numbers like ten out of 10,000 are very common, which is much lower than we would see, say, in airline markets.

Now, another feature is that because these sellers are small, they are frequently selling kind of one unit, and by one unit, I mean, say, a pair of tickets. So, one pair of substitute products or one unit of substitute products. So, in this case, we don't have the inventory incentive, and in a theoretical model, the declining price prediction arises very directly.

Okay, so what does the paper do? And obviously I'm not going to try to present all of this here, because there isn't the time, but I just want to give you an outline of kind of everything that's going on in the paper.

So, the first part is descriptive and shows, using data from StubHub and from eBay, both list and transaction prices, and different kinds of list prices fall by very significant amounts as the game approaches. And I'll just show you some pictures of that in a second. It's very statistically significant, see two statistics of 40 or more, but also very kind of economically significant.

Now, the next thing I do is describe various theories for why sellers cut prices over time, right? So, one theory is this kind of revenue management explanation, and just to be clear, what I mean by that is as the game approaches, your ability to sell tickets in the future goes down, you become keen as a seller, and therefore, you tend to cut prices.

Now, an alternative explanation is, let's say, residual man is becoming more elastic over time, right? So, maybe consumers who arrive near the end, they are going to have different kinds of demands than consumers who arrive early, and that would cause you to cut prices. And a third alternative is actually kind of a seller learning story, where because you don't have demand, you want to start off with a high price, learn about demand, and then cut prices sequentially.

So, what I do in the paper is I reject the seller learning explanation by testing kind of reduced form implications of a learn model, and then I use a structural model -- estimated structural model of the seller's pricing problem to distinguish between theories one and two and send up supporting the revenue management motivation.

Now, that analysis focused on what the seller -- on the seller's incentives, why the seller is cutting prices, and that's kind of consistent with the bulk of the theoretical revenue management literature. There is obviously the issue of why are some buyers willing to buy early if prices are expected to decline so much? Okay, and this has been recently addressed in the theoretical revenue management literature, and the kind of things that literature is focused on is what if you have risk aversion, if you have some uncertainty about availability, or you have search costs on the part of consumers, then maybe they'll be willing to buy earlier at higher prices.

So, what I do is I calibrate a utility function to look at the question of, okay, how large do search costs, say, or aversion have to be to justify early buying? And what I find is that you would expect, for example, if there's -- if there's return-to-market costs of, say, $25 per ticket, that would rationalize early purchasing. And then what you can do is you can look in the data, look at who buys early, and what you find is, for example, people who have to -- who are likely to have to make a lot of complementary investments to go to the game, such as people who live a long way away from the stadium, people who would, therefore, like to have return-to-market costs, they are the guys who tend to purchase earlier.

Similarly, people who are buying tickets for which future availability is more likely to be uncertain, they are also tending to purchase earlier. So, the patterns we see today are kind of consistent with the buyer side of these models, explanations as well.

Okay, so I just want to talk about two different things. Firstly, very briefly, provide some evidence on the fact that prices do decline and then explain what I do, separating the revenue management and residual demand explanations.

Okay, so a basic kind of linear estimating equation, where we are going to have price on the left-hand side, sort of dummies reflecting how many days there are to go until the game, list -- we have got listing characteristics, selling characteristics, some variables measuring kind of how the teams are doing. So, if a team suddenly gets in play-off contention, prices tend to increase. And then the sort of fixed effects, and the fixed effects are very important, because obviously you might be concerned that what's happening is that prices are falling over time because unobserved ticket quality is declining over time.

But, in fact, what I am going to have in my data is huge number of observations allowing me to include fixed effects, which I regard as being pretty exhaustive. So, for example, you know, very precise, you know, controlling for the actual row -- section and row the thing is in, including fixed effects for that, and then also looking kind of within seller -- within tickets, so the same seller selling the same ticket, what are they doing over time? And that's how I try and control the quality.

Okay. So, what are we seeing happen to prices? So, as I mentioned, I am going to be using data from StubHub, which is going to consist of data on list prices, and then eBay, where I have list prices for fixed-price listings, for auction listings, and buy-it-now listings, and also transaction prices. Here, we see in, say, the 40 days before the game, transaction prices on eBay falling by about 25 percent until the day the game is played and list prices on StubHub falling by around 30 percent, okay? And what you can see is the effects are estimated very precisely.

Now, when you look -- you can also see similar effects -- I should say larger effects when you look at e-Bay list prices and eBay auction start prices, and similarly, when you look within seller-ticket combination, you actually see, once again, even larger effects. So, for auction listings, you see the auction start price falling by 100 percent over time, and for list prices, you see declines of 60 or 70 percent kind of within ticket-seller combination.

Okay, so now let's turn to the question of why are sellers cutting prices? So, just to kind of fix ideas in a particular model, suppose there were two periods -- okay. So, suppose there were two periods, each period the seller has to set a price, unless, suppose, if after the second period the ticket is left unsold, he goes to value V, right? And that could be he actually goes to the game himself or he gives it away to a friend, for example. So, if you just set up the pricing problem, in the second period, the marginal -- what I'll call the opportunity cost of selling is equal to V. If you sell the ticket in the second period, the amount -- you're getting the revenue, but what you're forgoing is V.

In the first period, what you're forgoing is a complex combination of V and the revenue you would have got in the second period, right? So, what it's very easy to show is that if the demand function, this probability of sell function, Q, is the same in each period, prices will fall -- your prices should fall over time. On the other hand, there's -- and I'm going to call that the revenue management explanation.

On the other hand, there's -- prices could also be declining even if sellers are pricing in a very kind of myopic way, because the Q function is changing. So, if the Q function is becoming more elastic over time, that would also cause sellers to cut prices.

Okay, so, what am I going to do empirically? Well, at each point in time when the seller is listing his ticket, what he's doing is he's setting the price to maximize his revenues, recognizing that his cost of selling is this opportunity cost. The opportunity cost is not just V, but also reflects his future selling opportunities.

So, what I do is I estimate, in the case of fixed-price listings, a parameterized probability of sale function, allowing that slope of that and the intercept of that to vary over time, and then also I control it for, for example, the competition effects. I'm going to instrument the unobserved ticket quality using some instruments which seem to work reasonably well in this setting.

And then what am I going to do? I'm going -- from these estimates, I'm going to back out what's happening to the opportunity costs of sale over time, and then I'm going to perform counterfactual experiments which say, okay, let's suppose we take these opportunity costs as given, we hold fixed the demand effects and competition effects over time, how would that change the prices we see in the data?

Okay, so let's just look at fixed-price listings for a second. From a couple of different specifications, what you see quite clearly is that as the game approaches, the distribution of opportunity cost shifts leftwards and tends to decline, and that's true in a model way, which is based on kind of prices in a secondary market based on the face value of the ticket or using a low price model.

Now, how long do I have left? Okay.

Now, the next part of the analysis, which speaks more to the issue of why the prices are falling, consists of these counterfactual experiments. So, here you see, in the first two lines, what's happening to actual prices over time. So, on average, 41 -- more than 41 days before the game, the secondary market -- average secondary market fixed price is $70, when ten days before the game, it's $53, okay? And you could also see that it was holding to the median price.

Now, the counterfactual says, okay, we have estimated a decline in opportunity costs, and we're going to assume they continue to decline in that way, but we're now going to assume pay-off maximization on the part of the seller, but we're going to assume he's maximizing against the same demand curve in each period, and because competition also tends to increase over time, because there's more listings towards the end, we are also going to assume he faces the same competition as the game approaches.

What we see is we actually get price declines sort of very similar to those we actually see in the data, and that's true of both the models, and we see similar effects when we look at auction list -- different kinds of auction listings as well. So, what this tells us is the price-cutting patterns that we see are driven by the declining of opportunity costs and not by the changing elasticity of demand.

Okay, so I'll just wrap up, so -- allow more time for discussion. So, what I provide is very robust evidence for this tendency for prices to fall over time that hasn't really been established before in the literature in a perishable goods market. I then show, you know, in a number of different ways why -- a number of different kinds of pricing mechanisms, why sellers are cutting prices over time, and it's because of the motivation identified in the revenue management literature.

I then look at what's going on with early buying and find that the plausible levels of, for example, return-to-market costs, you can rationalize early buying, but there are some outstanding questions here. So, if we think airline markets or, say, advertising markets, so say radio and television advertising markets, we see different patterns there. Why is that the case? Is it because of differences in demand or, perhaps more intriguingly, it's because commitment plays a role in those markets, which it doesn't where we have tiny sellers in Major League Baseball secondary ticket markets?

On the other hand, also, what's driving the choice of selling mechanisms? So, you see very interesting shifts on eBay as the game approaches about what mechanisms people want to use, and this actually provides a nice environment, I think, for understanding what are the characteristics at different kind of sales mechanisms, which could have implications way beyond kind of this kind of market.

MR. STERN: And our discussant is Ian Gail, and let me pull up your thing.

MR. GALE: Okay. So, my introductory remarks have already been preempted, because I was going to start off by saying that the two big issues this week on The Washington Post have been, A, the election, and B, ticket resale. In fact, there was an article today where Michael Wilbon had initially talked himself to there being 15,000 Steeler fans at the game. By the end of his article, he was up over 25,000, and now we have got an estimate of apparently 45,000 if it was, indeed, half of the fans. StubHub actually claims it was -- they've sold 9000 tickets, but anyway . . . the problem is it's a big market, and it's a really, really growing market.

So, this is a -- this is a wonderful paper. I really enjoyed reading the paper. So, Andrew looks at two different markets, StubHub and the anonymous market, too, in this -- in this market, and a very robust finding is that prices decline over time. So, it's a substantial decline, but it's also not precipitous -- it's not a precipitous decline. It's fairly gradual over time.

Moreover, this decline shows up no matter how you cut the data. I mean, just stratify -- just look at all sorts of -- quality levels and so on, and the same pattern shows up.

All right. So, why declining prices? Well, actually, this is something I've thought about in the past, at least theoretically, and so there are different reasons why we got declining prices in the past. So, one reason would be if there was some market power that accrued if you were able to get the lion's share of the good in question. So, if you're selling off take-off and landing slots, for example, or if you had increased your returns to scale. In that case, winning the early -- winning the early units gave you a leg-up. You had a greater value of the later units than you did those who didn't have any units yet. So, that was one source of declining prices in auctions.

Another one, which is actually related, and it was actually cited, occurs when you have horizontally differentiated products being sold, because then, you may have a preferred product -- we could be talking about selling condominiums or we could be talking about selling baseball tickets -- you've got a preferred ticket, and simply buying earlier gives you the opportunity to choose which of the preferred -- which of the objects you get. All right.

Now, of course, we also saw rising prices, and so when we talked about airline pricing, naturally enough, people differed in terms of cost of locking in. That's something that Andrew talked about. So, those who had a lower cost of locking in would buy a ticket early for the -- for an airline, and, in fact, those people would be shunted to the off-peak flights, and those who showed up at the time of the flight would pay a higher fee. Okay, so, some of those same issues show up here.

All right. So, what do we find in the current paper? Well, clearly it's just an option value. I mean, you've got some value of going to the -- you could go to the game yourself, right? So, it's worth $50 to you to go to the game, but if you can sell it earlier for 70, you'll do that, okay? So, there's just a declining option value over time. As the clock keeps ticking, your chances to sell to somebody else get fewer and fewer, and that's why the prices are dropping. And empirically, this turns out to be the best explanation.

So, obviously, Andrew mentioned declining assets due to demand is another possibility or even learning by sellers could be. Sellers might not know this demand could be high for this game or this demand could be low, and that might lead to screening over time, okay, but that empirically is not what goes on here.

All right. So, then the question, of course, why buy early? Now, my first reaction had to be, well, how likely is it you're going to stock out? How likely is it that there are going to be no seats available, right? So, I can understand for Wrigley Field or Fenway Park, that's the case, but, of course, what's really important here is that it's not the case that you literally have no seats left. It's that you're looking for a particular kind of seat, right?

So, some of us are going looking for seats right behind the dugout and some of us are looking for seats in the bleachers, and so it may, indeed, be that the kind of ticket you're looking for may stock out even though there are many, many other seats left in the stadium, okay? So, that was the -- that was the important point.

In fact, one thing I didn't catch, but you may have had it in there, are there big differences across teams in the price paths, because that might tell you something. If, for example, the experience with Red Sox tickets or Cubs tickets differs from the experience with, say, Washington Nationals tickets. At least we know for the Nationals, the prices are going to be a lot lower.

UNKNOWN SPEAKER: They're very similar.

MR. GALE: Sorry?

UNKNOWN SPEAKER: They're very similar.

MR. GALE: They're very similar, okay. That's interesting.

All right. And, of course, other reasons why you are going to buy early, you know, search costs, and, of course, the complementary investments, which is something that shows up with attendance at a baseball game or a football game, but it also shows up in a newspaper a different form when we're talking about airline tickets.

All right. So, one question I had, you know, which tickets are available for resale? Is there any kind of selection issue? Now, the paper is very, very careful in controlling for -- controlling for quality, but a natural question would have been, is it the case that people decide to hold on to their tickets if it looks like the game is going to be good, but they dump their tickets if it looks like the game is going to be bad? Now, of course, you've controlled for all sorts of things, like how close the pennant race is, the standings, and so on. So, I guess I can't say much more than that, except you've obviously thought about that issue.

Who does the selling? So, another question that came to mind has to do with to what extent the teams are involved. Now, your data come from, what, 2000-2007, is that right? Teams have started to get a little more involved in selling tickets themselves. So, the Chicago Cubs now sell tickets through their Wrigley I premium -- it's called Wrigley I premium field service, whatever it is. Ticket service? All right.

So, basically the Cubs sell through this separate -- this separate arm, and, of course, the question is, why would they do that? There are many reasons why they might do that. One reason I've heard is that there's revenue-sharing in Major League Baseball, and if they can have their subsidiary making profits from ticket resale rather than getting themselves as primary sellers, that's to their advantage.

All right. I should also note, by the way, this past year, it's alleged that the Milwaukee Brewers sold a huge chunk of their tickets, playoff tickets, through a reseller, but the point is, that's only recent, and given your time period, that's probably not much of a -- much of an issue. Okay. So, I guess I -- okay, so I mentioned the Cubs and the Brewers.

All right. Forget about the first two points. The issue about price changes. Now, you noted that with StubHub, you're looking at listings, and if a listing comes off, then that's interpreted as a sale -- not necessarily, okay, because here's -- you've noted that someone might list a ticket on StubHub and then might change the price, and you're capturing -- you're capturing that, and the first one is not deemed to be a transaction, correct? Okay, good, because that would have been another possible source of problem, and you've caught that as well. All right, very good. Very good.

Okay, so just to finish up, one question is what are the -- what are the applications? What have we learned here? There are other industries obviously with perishable goods, and note that hotels and airlines, while they have one kind of pricing strategy themselves, they're often selling through other entities now, right? So, hotels are selling through consolidators. Airlines are selling through PriceLine.

So, the price dynamics we see, if you call American Airlines, may be very different from the prices -- from the dynamics you see when we actually look at PriceLine. And so, I'd be curious to see whether you're getting the exact same dynamics you've seen with baseball in PriceLine airline tickets, even though the airlines themselves want to commit to high prices for the people who just walk up to the gate.

And so I guess my time's up, so I'll just leave it at that. Obviously, welfare would be another issue, but again, let me say, it's a very, very nice paper.

MR. STERN: We have time for just a few questions.

AUDIENCE MEMBER: (Off microphone.)

MR. STERN: We're behind. So, I'm managing my interests versus keeping us on the time schedule.

Okay, and our last paper is by Steve Puller, along with some number of co-authors. By the way, I did remember, in fact, what the proposed title for this session was when we were organizing it. It is, in fact, you know, "That's the Ticket." So, here we go, more on tickets.

MR. PULLER: So, I will continue on the theme of talking about ticket pricing, in this case in the airline industry, and actually, it's been brought up a variety of times at least in the past paper what could be going on in airline pricing. So, it's actually a great setup to what we're going to be looking at. So, this is joint work with my colleague, Steve Wiggins, who is here, and Anirban Sengupta, who is at The Analysis Group, and what we want to do is understand better what's driving price dispersion in airlines.

So, we all know, at least through our own casual empiricism, there's a fair amount of price dispersion. So, if you flew to Washington yesterday, like I did, and you had polled your fellow passengers, at least on my flight you would find out that half of them had been in Oregon knocking on doors in the past two weeks, but you'd also find out that there's a fair amount of dispersion on the fares that they would pay, even if you found people who took exactly the same flights.

And so what we seek to do is to assess two broad categories of what might drive that dispersion. More formally, my discussant, Nancy Rose, has a paper or a book chapter recently with Severin Bornstein where they look at within-route price dispersion. So, same carrier, same route, and find a coefficient of variation in recent years and magnitude of about 0.5 or 0.6, and that's what we want to seek to understand, is within-carrier route price dispersion.

So, the two broad categories of theories that could describe that, one we're going to call scarcity pricing, and this is basically playing off some key features of the airline industry. First, the airlines have large fixed costs; the seat is perishable, as has been mentioned previously, so an empty seat takes off, loses value; and demand is uncertain.

The two primary theories that we're going to rely upon to develop comparative statics by Dana and by Gale and Holmes, and I'll give you a little description of that in a moment, but essentially the comparative static test which we are going to use turn on comparing features of tickets that are sold in high demand versus low demand states, or in the lingo of airlines, in the high load factor versus low load factor flights.

And then the second set of theories, which we're going to call yield management or sometimes called revenue management, essentially hinge on attaching restrictions to tickets. So, for example, a nonrefundable ticket or required Saturday night stay that create fencing devices that might allow the airlines to implement some form of second-degree price discrimination. And what these theories are going to predict is that fares are primarily associated with ticket characteristics in a way that's possibly independent of how full that flight actually is.

Now, clearly, these two theories are not mutually exclusive, and, you know, I'd certainly think that in airline pricing, what's going on is in combination of both of these, but what we'd like to do is try and understand which of these is the primary drivers of within carrier route price dispersion that we're seeing.

Now, one reason this has been difficult to assess in the past is one of the most commonly used data sets looking at airline pricing, data bank 1A. Before the Department of Transportation -- which is a 10 percent census of all domestic tickets. Before, the Department of Transportation releases it, they strip it of information which would be useful; in particular, they strip it of ticket characteristics. So, we don't know about refundability, and they strip it of information that allows you to pinpoint what particular flight that was. So, we can't assess load factor.

And so what we have is a unique data set which we hope will allow us -- which has load factoring and ticket characteristics, which hopefully will allow us to assess these characteristics.

So, just to give you a preview of findings, we're going to have some comparative static tests that compare characteristics of flights on high versus low load factor, and what we find is evidence that I think Steve and I would characterize as modest support of this -- of these scarcity pricing models, but it seems that there's stronger evidence that ticket characteristics are driving the dispersion that we're seeing.

So, just to give you a picture of the data, so you know what we're working with here, and I'll give you a little more description later. Essentially, we have a census of transactions through one of the major computer reservation systems. So, each data point here is a round-trip itinerary from Dallas-Fort Worth to Los Angeles International Airport in the fourth quarter of

2004, and we're graphing this versus the number of days in advance that the ticket was purchased, and each -- I guess it's red dot here is the mean per that day of -- for that day in advance.

So, as you can see, not surprisingly, as you wait until departure to purchase, in equilibrium, you're going to be paying a higher price, but there's a fair amount of dispersion around that, even as you come very close to the departure date, and it's that dispersion that we're seeking to explain.

So, let me describe briefly what the theoretical models are that are giving us these comparative statics. So, the first model by Jim Dana actually expands on earlier work by Prescott and Eden, and he gives what I think is a really intuitive example of stadium seating. So, I'll just describe that intuition.

So, imagine there's a competitive -- there's a stadium owner that is going to sell each of those seats competitively.  Prices are set in advance, which means there's a ticket price that -- a ticket printed with a price on it.  There is no resale in this model.  There is two demand states that will occur with equal probability.  And consumers vary in their willingness to pay.  And consumers show up randomly and they take whatever the cheapest ticket that's still available when they show up to buy.

So, imagine the marginal cost of capacity is 20. What Dana shows is that in equilibrium, firms are going to have intrafirm price dispersion. So, the stadium offer is going to offer a certain number of tickets priced at marginal cost. They are going to sell in both time -- in both high and low demand states, and then offer another quantity of tickets that's going to sell at a higher price, that's going to sell in only the high demand states. So, essentially, this is exploiting a zero profit condition, that the expected revenue from each ticket exactly covers the capacity cost.

Now, Dana expands upon this and gives a monopoly and oligopoly version. So, essentially what this tells us, the insight here is that one doesn't necessarily need fencing devices in order to get intrafirm price dispersion.

So, what predictions are going to come out of this? So, ideally, what the analyst would like to observe is multiple realizations of the same flight -- of flights that have the same expected load factor. So, if there's an expected -- the same expected load factor, the airline is going to offer different sets of tickets. Let's see if I can get this. So, some set of tickets that's going to sell with high probability, some that are going to sell with low probability. And what we'd like to see is multiple realizations, so that -- because our data are going to be transaction data.

So, in particular -- you know, so, for example, here's a possible realization, where 60-something percent of the tickets are sold. Sometimes, we're actually going to see less sold and sometimes we're going to see more sold. So, we're going to exploit that comparative static. So, what we want to see is for multiple realizations of flights with the same expected load factor.

On flights that have higher realized demand, you are going to have higher mean transacted fares, because you're kind of climbing further up the fare schedule. There's going to be more fare dispersion in transacted fares. There's going to be a larger share of high-priced tickets on the high load factor flights. And if you have data on the sequencing of purchases, you're going to find that for flights that are unusually full, say seven days in advance, the tickets sold in the last seven days are going to be higher priced. So, we're going to take these comparative statics to our data.

And the second model we're going to exploit is a model by Gale and Holmes, slightly different setting.

It's a monopoly airline. There are two flights, a peak flight and an off-peak flight. The consumers are going to prefer one of those two flights, but they don't know which one they prefer until right before departure time. So, imagine there's some business meeting, you don't exactly know when the meeting's going to occur until right beforehand, and these consumers are heterogenous in the sense that they vary in their time cost of waiting.

So, what the paper shows is that airlines are going to offer discounted advance purchase seats on the off-peak flight, and essentially what this is doing is diverting low time cost of waiting customers to the off peak flight. And so the prediction we're going to take to our data is to test if peak flights have fewer discounted advance purchase fares, where we're going to define advance purchase as those sales that were made two to four weeks before departure, okay? So, those are the two models within what we're calling scarcity pricing.

A separate literature in yield management, we're not going to directly test this, but just to make clear how we see this as a different type of model, this is basically playing off the fact that airlines might use ticketing restrictions like refundability to segment customers by their willingness to pay, and the key feature of these models in terms of predictions is different from the scarcity pricing models, is that the yield management models don't necessarily yield sharp predictions about the characteristics of tickets sold on high load factor versus low load factor flights.

Okay, so the data we have, I've made references a couple times, we have a census of transactions through the fourth quarter of '04 through one of the computer reservation systems that sells about a third of all tickets through all of the major distribution channels. So, for this, we have got more information than is in data bank 1A, because we know for each itinerary the various flights. We know which specific flight was taken. We -- for each coupon in the itinerary. And we know when it was purchased and what -- the days of the flights. And we can combine this with some other information to come up with a measure of what a given flight's load factor is, measured with a little error. And we're going to study 90 large routes for six large carriers, essentially the biggest carriers, with the exception of Southwest.

So, we also want information on ticket characteristics or restrictions. So, we went to another -- we went through another computer reservation system. It turns out that there's a historical archive so that one can look up what fares were. So, we're able to match -- gather information from that second computer reservation system on ticket characteristics.

The key one is refundability, where there's a restriction on the days one could travel, and whether there's minimum or maximum stay restriction, typically like a minimum one-day stay restriction. So, we matched that to our transaction data. We were able to match about 36 percent of observations. We do some tests which are in the papers where we're comparing characteristics of the matched and unmatched, and in our view, they're not that different. So, we think the matched transactions are reasonably representative.

So, before I go into the formal test, let me show you some motivating regressions that we view as kind of descriptive analysis of the data, which are then going to be consistent with some of our formal tests. So, what we do in these regressions is we're just regressing logged fares on ticket characteristics and then on load factor.

So, in the first regression, we're only using ticket characteristics. As you can see, as you buy closer to departure date, you are going to be paying more. A refundable ticket has about a 50 percent premium. Travel and stay restrictions, so, for example, a minimum stay restriction, corresponds to paying 8 percent less. And tickets where the traveler stays over a Saturday night is a -- are transacted at a 13 percent lower price. And from there, we're explaining about 70 percent of the variation.

Now we're going to add in load factor, both actual load factor and expected load factor, and see what the signs of those coefficients are. So, when we add actual load factor, it is statistically significant but economically fairly small, a standard deviation increase in actual load factor means the ticket will sell at 1.5 percent higher fare, and as you notice, the amount of variation that we explain is not a lot higher.

We add in a measure of expected load factor, and you also get that for flights expected to be more full, that the fares are slightly higher. And you add both of them in, and it seems that it's expected load factor that is driving this.

So, from this we take, I guess, two things. First, to the extent that load factor impacts fares and equilibrium, it is economically rather modest; and secondly, if you look at the coefficients of the ticket characteristics after we add load factor, the coefficients are fairly robust, which suggests that to the extent that load factors is impacting fares, it does so in a manner that's largely independent of the load factor.

Okay. So, now to our formal test. So, the first test is that for flights with the same expected load factor -- if they have a higher realized load factor, you are going to see higher price dispersion. So, how do we measure this? So, what we do is we take a given flight on a given day of the week. So, say American flight 301 on Monday. We see that happen 12 times in our sample. We are going to take the average of that as our expected load factor, and then we divide all flights into quartiles.

Then, within each quartile of expected load factor, we take the realized load factor, again divided into quartiles. So, what the theory would predict is that once I fill in these cells with some measure of dispersion, that for a given column, a given expected load factor, as you move up the column, fuller flights, you're going to have a higher measure of dispersion.

So, the dispersion we're going to use is the Gini coefficient. So, let me fill in those cells there. So, as you look, as you move up any of those columns, you're actually finding that, if anything, there's actually a decrease in dispersion, although it's not economically that large. So, it doesn't appear that there's substantially more dispersion on flights that are ex post full.

The second prediction is we're going to test how different types of tickets -- quantities of tickets that are allocated between different categories. So, the main predictions that come out of the two theories are that the share of low price tickets is going to be lower in high demand states, and off -- on-peak flights are going to have a smaller share of advance purchase or discount sales. So, to do this, we're going to have to -- we're going to have to stick tickets into categories and define exactly what a discount ticket is.

So, the way that we've approached this -- and we're welcome to other suggestions -- is to group tickets into three different groups, basically high, medium, and low. So, group one is going to be all of our refundable tickets. These tend to be tickets that do not have restrictions. Group two is going to be nonrefundable tickets that have no restrictions on travel or stay. And the residual category, group three, is going to be nonrefundable tickets that have a travel and/or stay restriction. And as you can see, these are largely consistent with prices. The relative prices are higher in group one versus group three.

Okay, so now to test these, I'm going to show you a series of tables that look like this. So, let me make sure I'm clear on -- clear to you guys on what exactly these tables are. So, what we want to do is we want to look at flights that are off peak and peak. So, for an off-peak flight, we're going to take flights that were expected to be low load factor and are load factor, so no unusual shocks. Peak, expected high, realized to be high. So, on these tables I show you, the left-hand side is going to be our measure of the off -- allocation for the off-peak flights. The right-hand side is going to be the peak flights. And the boxes here are basically percentages of all tickets sold for each of these various carriers.

So, we're testing here, at least in the Dana theory, is whether the fraction of seats sold on off-peak flights that are low, discount tickets, group three, is higher than the sale of tickets on peak flights. As you can see, for American, it's 53 versus 50, so, in fact, the sign is right, although the magnitude seems relatively modest. For Delta, it's 55 versus 49; for Continental, 53 versus 46; for other six carriers, you can see that there is a modest reallocation, it appears, but in percentage terms, it's, you know, between like 2 to 7 percent of the tickets.

Now, let's directly test the Gale and Holmes prediction about advance purchases. So, we're going to define an advance purchase low-priced ticket, pulling from the Gale and Holmes theory, as a group three ticket that is sold two weeks or more before departure. So, again, we're comparing left, off peak, to the right, peak, and you can see the difference is, across the airlines, there's zero percent difference, 3 percent difference, 4 percent difference, for other carriers, 4, 2, and 4 percent. So, again, it is consistent with the theory, but it doesn't seem that there's a large fraction of tickets that's actually reallocated to the -- low-priced tickets as reallocated to the off-peak flights.

How much time do I have? Two minutes, okay.

So, let me go through our last test. So, the last test, we want to ask the question, since we observed -- we want to exploit the fact that we see the sequencing of purchases. So, we actually know how many tickets were purchased as of seven days before, six days before, five days before departure. So, what we want to ask is, for a flight that is unusually full seven days before, are the tickets that are sold in the last seven days unusually higher priced? And that would be consistent with climbing up a fixed fare schedule, climbing up that fare schedule that I had drawn before.

So, let me leave the -- let me just kind of focus on the bottom question here, in yellow. So, essentially we're asking the question, for a ticket bought seven days before departure, if the plane is 10 percent fuller than normal -- and the paper describes exactly what normal is -- what percent more expensive is the fare?

So, what we did here is we put everything into -- we did a kernel regression for each of the carriers and estimated this effect separately by carrier. So, the steepest there is American, so that would suggest 10 percent fuller in the last seven days. The last seven days' tickets are priced 1.7 percent higher. For the other carriers, it seems to be smaller than that.

So, bringing all these tests together, we think that -- we interpret these various tests as kind of painting a similar picture, suggesting that, yes, in fact, there is evidence consistent with the scarcity pricing predictions, but it seems to be relatively economically modest in terms of the quantity of seats reallocated and what the pricing effects of that are.

It appears, in contrast, that there's much stronger evidence that ticket prices -- so think back to those motivating regressions -- that those ticket prices are explaining more the variation. So, while it certainly doesn't rule out models based on what we're calling scarcity pricing, it certainly suggests that, at least in future research, it would be very interesting to look at how these ticket characteristics can be used to segment customers.

Thanks.

MR. STERN: Great. And our final discussant is Nancy Rose.

MS. ROSE: Okay. So, I wanted to say I found this paper extremely interesting, and I'll have a few comments of ways I think that the authors might push a little bit more on their data, but I think this is a fascinating insight into airline pricing that we really haven't been able to get before because of the lack of data that -- lack of information in the data sets we traditionally use.

So, let me first -- first start with a little bit of disclosure, which is I start from an extremely strong prior, that it would be very difficult to explain airline pricing without reference to at least some price discrimination, that while I also think you've got to refer to stochastic demand management, I think it's also impossible to understand airline pricing without thinking about that, that despite the ongoing debate in the economics literature over whether you need to resort to price discrimination to understand patterns, my prior is you do, and let me just be up front about that and mention a few reasons why I think that's true.

First of all, we've got the internal documents that we heard about often influencing judges, maybe not just judges, but economists as well, and so I had the treat early in my academic career to read a lengthy document from the internal pricing manager at Northwest Airlines that went on and on at some great length about the desire to find restrictions that were very costly for business travelers but not very costly for leisure travelers as a way of separating those two markets, referring to Saturday night stays as the single best restriction of them all, and referring to this as -- inexplicitly in price discrimination terms.

The second is over the course of the last decade, watching airlines become increasingly focused on not just setting ticket restrictions to segment those customer bases, but to find ways to make the restrictions more and more and more costly to their high willingness to pay travelers, and so I think about the imposition of nonrefundability; the increased fees for flying stand-by or day-of-flight changes; even the efforts by some airlines to go after their most high mileage frequent flyer members and revoking miles because those passengers took liberties with tickets that weren't in accord with the contract terms of carriage.

So, unlike the first two papers in this session, airlines put an enormous amount of effort into preventing resale, and as I teach my undergraduates, for price discrimination, you need market power, some segmentation mechanism, and thirdly and crucially, an ability to prevent arbitrage or resale.

And it just seems to me the fact that airlines are putting so much effort into this, and at times helped, for instance, most recently by the U.S. Government deciding that for security reasons, they would help to enforce the no resale option by requiring you to have a government ID that matched the name on your ticket. It just -- all of this gives me a strong suggestion -- strongly suggests to me that price discrimination is at work.

And I'll give you the last sort of kind of example that I give my undergraduates when we talk about airline pricing to suggest that there might be more than realized demand going on in terms of thinking about airline prices.

So, the first line here is the one-way fare for a ticket from Boston to Detroit on Northwest Airlines on an October 20th date, purchased or explored for purchase three weeks in advance, and looking for a one-way itinerary, just this flight, Northwest 371, and the ticket on that or the price of that ticket was $811, all right? So, maybe that was an extremely full flight. And if you added a return on October 22nd, two days later, maybe that helped, because that flight was really low load factor, but it's hard to reconcile a negative price with adding that second return flight to your itinerary as being explained by anything like stochastic demand management, right? We need so much to divert you to that return flight, we'll pay you $400 to take it or $350 to take it.

So, what this paper tries to do is to pay some -- take some care in constructing the comparative static implications of models of stochastic peak load pricing and then to also equally carefully construct data on ticket restrictions and mapping those into fares and try to give us some sense for how much of the variation in prices is left over once you look at these comparative static variations.

And what I want to emphasize to you that I think is fabulous about this paper -- and I expect it to be only the first in a series of papers from these authors -- is that they've put together an amazing data set that's got a third of all of the tickets sold in the U.S. in the fourth quarter of 2004 with very detailed information on those ticket characteristics, and I think just the descriptive statistics are fascinating.

So, I looked at the Table 6 that Steve gave you a little taste of, and, you know, just lots of interesting patterns jump out there that have nothing to do with either stochastic demand management or price discrimination, or not immediately. For instance, the large fraction of tickets that are sold in the last six days before flight departure, which, you know, we really didn't have aggregate data on that that would let us look at that, and I imagine you could do a lot with those tickets; or if you go -- flip through that at the break, the very high fraction of group one -- remember, those are the really high-priced, unrestricted tickets -- that U.S. Air sells, unlike all the other carriers.

My guess is -- and this might be something you guys want to do -- if you pulled the shuttle routes out of that, it might look a little different, but I thought that was fascinating, and clearly, having a high fraction of group one tickets isn't enough to make you a highly successful, profitable airline.  But anyway, the descriptive stats suggest a lot of interesting data to be mined yet in this project.

So, let me say, in the spirit of comments, just a few suggestions for things that we might do, and I'll give you my bottom line in case I run out of time and Scott cuts me off, which is this is a very interesting and I think innovative paper, and you definitely want to read it.  So, just a few suggestions.

The first is I'm a little uneasy with the implicit identification of high demand flights as those on which we observe high load factors, in part because I think if the revenue management systems were really perfect and -- or ticket restrictions were perfect and the revenue management systems were ideal, that the airline's goal is to maximize revenue, which isn't the same as having all the flights depart full, but it would push in that direction.

So, you'd adjust your allocations and your price engine a way to basically, if last-minute diversion wasn't an issue -- that is to say, if I priced that last-minute ticket at zero, I'd have a lot of people waiting and showing up to buy it at zero -- if that wasn't an issue, then airlines would be willing to sell any empty seat at zero.  The load factor would be a hundred. You couldn't tell anything about the underlying demand for a given flight.

So, I might just say a little more thought about that, is there a way to get a better indication of what demand is? And my thought was it might at least be interesting to look at the periods that are excluded from this analysis, which is the -- so, this is the fourth quarter, includes holiday travel. We know and the airlines know that every flight that flies out on the Tuesday or Wednesday before Thanksgiving is going to be full.

Now, the kind of demanders are different, so you might expect the price level to be different, but there's no probability of -- at typical airline prices, right, there's no probability of a seat going -- going out unsold on the Wednesday before Thanksgiving, apart from kind of price discrimination stories, maybe. That might let you do something with these models that predict that prices should be inversely proportional to the probability of selling the seat. So, maybe something that could be gleaned from looking at those periods.

The second -- and this is really more minor but I think might help to crispen some of -- or at least understanding it might help to crispen some of the results, which is I suspect, given that all of these routes are basically hub routes, the way that the routes are selected are the very largest routes, those are routes on which the hub carrier is likely to have corporate discount deals with significant employers in the area, and I'm guessing first that that explains why some of the fares don't perfectly match the restrictions, although you might be able to get within the range.

But secondly, that some of those -- those tickets that look like unrestricted, very low-priced tickets sold right before departure may, in fact, be corporate discount tickets that are -- I'm sorry, that look like restricted tickets sold right before departure are, in fact, unrestricted corporate discount tickets. So, you know, I bought my ticket two days ago to fly on the shuttle. I paid 230 because of a corporate discount rate. The unrestricted walk-up fare is 448. My ticket looks more like a discounted, advance purchase ticket, but it doesn't have restrictions associated with it. And you might be able to do something by looking at nonhub carriers out of hubs.

And then finally, I'd like to push them a little bit harder to say something more -- to take more seriously the revenue management literature and what the implications of stochastic peak load pricing also are maybe for expected revenue per seat or expected revenue per flight, and I'm wondering if you could do a little something more with that, again, to get crisper predictions that you could then test against the patterns and the data.

But overall, let me just say, I found this a fascinating paper, a great contribution to the literature, and I look forward to not just the revisions of this one, but the many papers that are to come.

Thanks.

MR. STERN: Okay. So, we have an incentive conflict, because if you have -- we probably have time for one or two questions, and then we have lunch, and so are there one or two questions that we want to do before?

And why don't --

AUDIENCE MEMBER: I have a question for all three of you. It seems to me that all of you have made assumption of exogenous entry of consumers. For Alan, I think it's random participation in the secondary market; for Andrew, it's also consumers arrive exogenously. I think you also mentioned that consumers arrive exogenously -- randomly. So, I'm wondering the impact of this assumption on your work.

For example, in Alan's work, if consumers -- I would imagine consumers who participate in the secondary market would be -- will have higher value of time, and I'm wondering maybe like you leave out very small welfare impact of the resale market is partly determined by that. And I think this assumption may have different impacts on all three of your work. So, I just want to know.

(Inaudible response.)

AUDIENCE MEMBER: So, what prevents you from making alternative assumption?

AUDIENCE MEMBER: (Inaudible response.)

UNKNOWN SPEAKER: Could you go to the microphone?

MR. PULLER: So, I think undoubtedly that the types of customers that arrive right before departure are different than the customers arriving maybe 30 days before. We're kind of taking the theory seriously as to what the predictions would be if it's purely scarcity pricing, but I would completely agree with you that it would be interesting to try and understand the characteristics of the customers that are arriving early versus late.

MR. SWEETING: (Off microphone.)

UNKNOWN SPEAKER: You can also look at time slots, which we have done some preliminary work on time slots. It's a way of getting the pure exogeneity of load factor, so you look at time slots and use that as an instrument, because the demand would be very systematically different from, say, 7:00 to 10:00 a.m. in the morning -- (off microphone).

MR. VITA: Any of you guys who are talking, you have to talk into the microphone.

UNKNOWN SPEAKER: But I think we're done.

MR. VITA: Okay, we're done.

(Paper Session One concluded.)

**KEYNOTE ADDRESS BY CARL SHAPIRO**

MR. SCHMIDT: Okay, it seems quiet enough that we can go ahead and get started. I just wanted to quickly say, I'm Dave Schmidt. I run our Office of Applied Research here in the Bureau of Economics. And I just wanted to send an invitation out to a lot of the researchers here that we think there are a lot of great complementarities between economics, doing research and people here at the agency, specifically at the FTC.

And if you've got like good empirical models or theoretical models that you're sort of looking for an application for, I encourage you to come and grab one of us and talk to us, because we have a lot of experience in industries and we might be able to think of something that fits well. I think there are great complementarities there between us and I think our lunchtime speaker here really embodies that.

He's recognized those complementarities through, I guess, integration of some sort by serving at the Department of Justice as the Deputy Assistant AG for Antitrust, and currently as the Transamerica Professor of Business Strategy at UC Berkeley. And we've all benefitted from Carl's research and writings and, now, today, from his lunchtime speech here on market definition.

So, Carl Shapiro.

**(Applause.)**

MR. SHAPIRO: Hi, let me get the slides set up here.

**(Brief pause.)**

MR. SHAPIRO: It's really a pleasure to be here. So, this is joint work with Joe Farrell. I actually spoke about an early version of these musings back in February at the Unilateral Effects Workshop in this very room, and we've now been floating these ideas a little bit.

It's interesting. This is not a talk about market definition; this is a talk about not market definition. So, the Gestalt here is, look, particularly among economists, years, if not decades, of wringing our hands about market definition is kind of -- I don't want to quite say screwy, but can be an indirect and confusing way to address the problems, the issues in certain mergers, particularly unilateral effects mergers and differentiated products. So, we are coming -- the idea here is to put forward a very specific alternative to going through the market definition market share exercise that we think is very practical, but has very good economic pedigree, as I will argue.

So, I'm going to be very quick breezing through current practice, assuming that most, if not all, of you are quite familiar with current merger practice under the horizontal merger guidelines.  And, of course, this does reflect case law as well.  So, I will turn to more of the legal issues a little bit, although it really is not my focus.

So, we define the relevant market.  There's a specific algorithm for doing that, including the snip test.  Then the point of that is to get a concentration, Herfindahl Index, measures of concentration, and the point of all that is so that the -- I will be presenting this somewhat from the perspective of a government agency investigating a merger and deciding whether to challenge. The point of that is for the government to see if they can obtain this structural presumption that the merger's anticompetitive, which then would have to be rebutted by the merging parties if they went to court.

I put a little question mark around the structural presumption because it ain't so strong anymore.  In the paper John Baker and I published -- actually, a couple papers in the last year -- goes through some of the history and explains what's happened to the structural presumption.  This is one of the issues is that because that structural presumption gets less weight, this whole exercise, this rather elaborate exercise, isn't giving as much pay-off anyhow.

And then we have the back -- what I'll call back end where this presumption could be rebutted by the parties through arguments about entry, repositioning, efficiencies. Or not just rebutted by the parties, but the agency itself, of course, could say, oh, we're not going to bring a case, even though we've got high concentration, because we think these other effects are strong.

When you're using a consumer welfare standard, there's a whole separate debate about whether that's the right standard, but that seems well-established and we're just totally accepting and embracing that in this proposal.

So, the goal here is actually, in some ways, to be rather modest, in some ways to be a bigger change in slotting into this structure an alternative way to get a presumption. Not to replace this method, but to offer a second root for the government to get a presumption, and then, continue on into the back end to see if we've rebutted in the same fashion. So, that's the idea, without using market definition.

Now, I'm really thinking of differentiated product mergers, of which there are plenty. I mean, Mike Vita and his whole shop, it sounds like they spend their

-- it’s a lot of what they do, actually, at least the ones you mentioned maybe in the context of merger simulation. And, you know, many of you who work on mergers will be -- you know, there’s just a wide class of mergers, okay?

Maybe my experience is a little -- not totally representative because I haven’t done steel mergers or chemical mergers very much, but a lot of the mergers, I think, would fit here, you know, that were differentiated products. And I listed some of the industries here. I didn’t even mention, you know, retailing tends to fit in here, too, or something like a hospital merger would be similar where the differentiation may be geographical as well.

So, a very large class, but we’re not claiming all mergers. So, it’s unilateral effects and these differentiated products.

We’re going to focus on pricing competition in the same way that the merger guidelines do in terms of thinking about measuring things and effects, which, as we understand, price is just one dimension of competition, but it’s a good sort of proxy for other dimensions, and when we try to measure effects, we’ll look at prices first and then -- at least initially. There’s some stuff in the paper about how you would apply these same ideas to innovation competition as well.

So, this class of mergers, as I said, is very large. And some of the most visible litigated mergers, obviously, just the tip of the iceberg, are in this area. Whole Foods -- in recent years, Whole Foods, certainly. Over at DOJ, the Oracle and Sunguard cases that -- both of which the Division lost, were exactly in this area. So, very important.

So, the alternative method that I wanted to develop, not meant to supplant the current approach. As I already indicated, the presumption that would be established would be rebuttable. So, when I talk about this whole method, you should not think about it, oh, that's the complete merger analysis any more than you should think that calculating the Herfindahls is the complete merger analysis. So, just think of it as substituting for that part.

But the idea is very similar. We're trying to come up with a simple rule which is, essentially, the first stage, and if the government can show certain things then the burdens shift.

In my paper with Jonathan Baker, and many other people have written, how really important it is to have some burden-shifting method because if the government really had to prove to some standard what the effects would be, it’s very hard, because these predictions are very hard. I mean, look, we have trouble retrospective, figuring out what happened even after the fact. So, I’m in the camp, certainly, and I know Joe agrees with me, that some ability to get a presumption is important if one is going to have effective merger control as a practical matter. That’s debatable, obviously, but that’s where we’re coming from.

So, part of the motivation here, I guess a large part of the motivation is that we, and I think many other economists and lawyers, see a lot of difficulties, problems, even dysfunction with the current system, this current structure based on market definition and concentration. I think the root of the problem is that that structure was put in place -- and it dovetails much better with coordinated effects cases than unilateral effects cases.

If you think about a case where you say, gee, these companies are merging, I’m concerned this is going to make it more likely some cartel -- I’ll just use that hard-edged word for coordination. A cartel will form.

You’d say, well, gee, maybe I should figure out which firms would have to be in the cartel and then figure out is this going to be significant change in that group of firms. So, if there are five firms that need to be in a cartel after this and there used to be six, I can think about it that way. And that's what this hypothetical monopolist test does. It tries to figure out what that group of firms is that would profitably run a cartel.

So, it's actually -- the market definition actually is well-matched. I mean, it still has difficulties of various sorts, but at least it's logically matched with the coordinated effects theories. And that's just not true for unilateral effects, okay? It's just not, okay? Because unilateral effects is really what's going on between these two firms given some cloud of other firms around them and drawing a boundary on which other set of firms we want to include or exclude becomes somewhat artificial and is not the direct question at hand.

And I make these assertions here. This method can be misleading, uninformative, distracting. Those of you who work in this area will, I think, mostly nod your head and agree. Those of you who aren't convinced, you will have to talk to me later or talk to other people who work in this area and, perhaps, you'll then be convinced.

It's also the case that the method, as described in the guidelines, introduces various, I say here, arbitrary parameters. There's a bunch of things that one has to do so you have the size of the snip. So, it’s 5 percent most of the time, but could be more or less, the HHI’s thresholds, 1,800, Delta, HHI of 100. Thirty-five percent safe harbor. These are all numbers that -- I don’t mean totally arbitrary like they were picked out at random, but, you know, where do they really come from? You know, what’s the basis for those and what’s the logic behind why it is 1,800 or a change in the Herfindahl of 100 or 150?

Maybe you could go back to some study of cross-sectional margins or something from the seventies, but it’s not going to work very well with these cases. It doesn’t really fit. So, you’re kind of -- you’re just plugging in these things and you can’t really explain it very well. At least I haven’t seen anybody explain it very well.

And this leads to all sorts of practical problems because judges kind of scratch their heads and they’re like, what am I supposed to do here? In Oracle, people saw the DOJ loss because they were unable to establish the relevant market in that case. There’s this critical loss methodology that’s been abused, and there’s nothing wrong with it at the level of arithmetic, but it’s causing problems because of the abuse.

And, as I said before, the pay-off to the whole exercise is reduced because what do we really make of it when we say, okay, the Herfindahl goes up from 2,200 to 2,700. Then the judge will say, oh, that's interesting, but does that tell me much? So, there you are. So, that's not establishing a very strong presumption except, I guess, in the case where it's very, very high concentration.

So, we want to come up with an alternative simple test diagnostic and the -- I want to draw a parallel here between the existing concentration base test and our test. UPP here stands for upward pricing pressure that I'll be talking about quite a bit in the slides to come.

So, if you think about, again, sort of the Gestalt of the Herfindahls, there is an underlying robust economic idea there, which is if you have a large share of the market and you increase output or lower price, you lose some revenues on all your inter-marginal units, as you have larger share. You may pick up some business from others. But as one firm becomes larger in the market, they have less incentive to increase output or lower price. It's just -- you can see that in the Cournot Model most clearly where Herfindahl type numbers can come up.

In my paper with Joe Farrell, actually, in the

AER and the other paper in the RAND Journal in '89 and '90, we went through -- and other people have looked at the relationship between Herfindahls and Cournot. But there is an underlying robust idea relating share to marginal revenue. It really works pretty well. I mean, that idea really fits more with homogenous products and output choices and not so well with differentiated products.

So, what we're doing is saying, well, let's look at -- think more at this level, the Gestalt level of differentiated product and pricing, and we're going to look to see whether the merger will create pressure for prices to go up, a very robust idea that we will develop in a specific model of Bertrand competition, but the idea will be very robust, just as concentration-based ideas are robust, although they're developed in a rather specific model of Cournot. That's the parallel.

Now, let's develop that test. That next group of slides does that. So, in a way what you might do is draw a line here, take a fresh start. Put the guidelines aside, as difficult as that might be for some of you. Just imagine you're taking a clean sheet approach to thinking about how you would evaluate a merger, differentiate a product industry, what would you do? And you'd say, I think two things. You'd say, well, gee, these companies have been competing against each other, also against others. But what the merger's going to do is they're going to stop competing against each other. That's what we mean by unilateral effects. And that will generally encourage some higher prices. That's a pretty general idea.

On the other hand, they do get to combine -- there may well be some efficiencies by combining various assets. If that's successful and it lowers costs, maybe marginal costs, maybe other costs, that will tend to push towards lower prices. So, we've got to have a trade-off between these two.

So, we're going to try to look at whether that trade-off will tend to push prices up or down. And it's very important, and you'll see the trick here, if you will, is that we're going to look at the direction of the price change, but not worry ourselves at this presumption, at this point where we're trying to establish simple rules of presumption, we're not going to worry ourselves over trying to predict the magnitude of the price change. That's the key simplification that gives us mileage. I'll spend quite a bit of time talking about that in a moment.

So, that's the basic trade-off, loss of direct competition, but efficiencies. So, we're going to build that right in. We're just going to go right there. We're going to go right to that without any -- you know, nobody's telling us some other artificial construct. We're just going to do that directly, okay?

So, what does that mean? Well, the loss of direct competition, we can -- I'd use the term "cannibalization." Of course, what I mean by that is before we merged, when I got business from you, that was in addition to my product, after I acquire your product, if I manage my product to get more sales from your product, I'm now cannibalizing my own sales. So, that's not nearly as big a win as it was before the merger.

So, how do we think about that? So, we've got some notations. We've got two firms, let's say. The very simplest structure you could think about this: Two firms, profit levels. First, I'll do it abstractly and then in the next slide I'll talk about prices. Just think generally that I run product -- I'm Firm A. There's some strategy beyond variable. If I do more on that dimension, I sell more. X is output. But if I do that, it cuts into your profits. This might be lowering prices. It might be marketing more. It could be improving my product. I don't care.

So, the merger internalizes this impact. That's what I mean by cannibalization. So, what you can figure out is, if I think about the -- if I think of the cost of selling one more unit of my product, well, this cannibalization is equivalent to a cost increase, I mean, marginal cost increase there for my product. That's what I call Tax A. And it's basically -- which is how much your profits fall if I sell one more unit. That's basically an opportunity cost term that gets internalized through the merger.

And that's what that ratio is, how much your profits fall by -- if I increase my -- if I act more competitively, how much your profits fall per unit extra that I sell. We could think of trying to measure this. In the abstract, obviously, that's pretty hard.

So, let's talk about pricing. We'll talk about pricing. So, now, if my strategy variable is price, so we've set prices. We have marginal costs there. This, of course, now brings in the diversion ratio, which is just a -- I think, for many people, an easier-to-understand version of cross-price elasticities.

So, diversion ratio is if I sell one more unit, what fraction of units do you lose? So, if we're pretty close competitors, maybe that would be a half. Whatever I do, half my sales come from you, half from other firms, other products. Obviously, it could be a lot lower. Maybe you think half is a lot. That's the number.

So, we measure that and talked a lot about that. Of course, that came up earlier this morning. So, now, we figure, well, how do we then come up with this opportunity cost term or tax? Well, if I sell one more unit that comes at your expense, based on the diversion ratio. And then how much -- if you lost sales, how much does that matter? Well, that depends on your margin. That's the effect on your profits. It would be how many sales you lose and how much your profit margin was on those sales. So, P2 minus C2, P is your price, C is your marginal cost.

So, this is the opportunity cost term. So, if I sell one more unit, that's the cost in terms of profits on your product, which is now internalized due to the merger. So, that's one side of the equation. That's the loss of competition. So, we actually then kind of have a way to quantify that.

Now, what about the other side of it? Now, here's the difference. In the merger guidelines, there would be this whole work-up and then only at the end do we say, oh, we figured out all this stuff, now let's compare efficiencies. So, one of the other tricks here is we say, no, no, no, what we should be doing is we've got to -- basically, the merger we can think of as like a cost increase for the product I'm selling. It's an opportunity cost cannibalization term. Again, I just showed you the expression for that. So, I've got a cost increase on product one that's D(1)(2) times P2 minus D2.

Well, is there an offsetting cost decrease? Let's compare those right now at the front of the analysis. That's the ultimate trade-off in merger policy anyhow. So, what is the cost decrease that might counteract that? And we're focusing on my product, Product 1, at this moment. We can do this for each product. Well, the margin of cost is C1, let's assume that there's some efficiency that we're willing to give automatic credit to the merger, some fraction of my marginal cost. And that fraction, we're using the term "E" here for the efficiencies.

This is what Rick Warren-Boulton 20 years ago called the standard deduction. We're going to credit the merger with some efficiencies without requiring the parties to show it. In the back end, maybe they could show the efficiencies were greater or maybe the government could show they were less, but this is what we're willing to credit mergers with. This would be a policy parameter.

Presumably, I mean, based on general evidence about how mergers actually do lower costs and you could imagine, if you wanted to get sophisticated, tuning it by industry or type of merger or something. But that's a level of sophistication I don't want to get into, just the same way we don't have different Herfindahl thresholds for different industries.

So, again, don't think of this as a complete analysis, think of it as supposed to be uniform, simple, transparent so that companies can rely on this. So, we've got a policy parameter, E.

One of the differences here, though, as you can see is, look, at least we would know what empirical evidence to look at to figure out what E should be. I guess you would know what to look at to figure out the Herfindahl 1,800 and the change of 200. I haven't seen a lot of studies that really tie that down for me, why that's 1,800 and the change is 100.

Here, I'm not saying it's going to be easy, but this is what you really care about. What has actually happened when firms have merged? What sort of efficiencies have they achieved? That would be built into the structure.

So, now we can put the two pieces together. We're still on our clean sheet. It's only half full so far, but it's still pristine, a rather nice, crisp piece of paper. So, we ask which of these forces is stronger? So, if you just think about my product, I had this cost increase term which was the cannibalization effect. I have a cost decrease term, which is the automatic credit for efficiencies. So, which one's bigger?

Well, there it is. The left-hand side is the cannibalization term; the right-hand side is the efficiencies term. So, we would say we have upward pricing pressure for product one if this inequality is satisfied.

Now, I snuck in putting bars on the variables here just to indicate that what -- to make this practical, what we're going to do is measure prices and costs at pre-merger levels. So, that's what the bars mean. And I'll talk in a moment about how that biases the test results in certain ways that I'll get to. Basically, it leads to false negatives and I'll explain.

So, this is the key inequality. So, you now can take out your guidelines, you can circle the parts that have the Herfindahls and you say, Option B, you put this. I did not tell you to strike them out and erase those parts, but here's Option B. Now, that's the argument.

In a case where there are -- if you had a symmetric case with the two products being symmetric, you can write this a little differently in terms of the price cost margin. You can simplify it. Just a couple lines of algebra. But this, I think, is helpful for getting a sense of the calibration of the test.

So, for example, if you had a margin of one-third, not an unreasonable margin in a lot of industries -- it could be much higher in some industries -- if you were willing to spot 10 percent as your efficiencies associated with the merger, then, in that case, you would get upper pricing pressure if the diversion ratio was greater than 20 percent. You got a symmetric situation. You can put in other numbers. The paper has more formulas and examples.

So, now, we have a theorem. Even though it's a policy paper, we have theorem. So, I'm going to now say if this inequality is satisfied, I'll say there's upper pricing pressure for Product 1. Okay, let's define that. The theorem says, if there's upper pricing pressure for both products and they merged and the merger caused the default or assumed level of efficiencies, both prices would go up in a Bertrand duopoly. So, that's the specific model that sort of underlies the logic here.

Now, I know very well, and you do, too, and we heard earlier today, not all these margins are going to be Bertrand duopolies. But this is -- and I will talk about that, too. But this is the simple logic underlying this in theorem form.

The concept, I think, is very robust. You know, loss of direct competition compared with efficiencies, but to actually come up with an operational test, one does need to make some additional assumptions and we've got that here.

Now, this does not say how much the price is going to go up. I didn't say anything about pre and post-merger equilibrium. I just said the price will be higher. So, that's the result. So, that's where we're going.

Now, let me detour a little bit and say -- some of the objections we've gotten to this over the past half-year or so when we've been kicking it around. Some people say, well, okay, you've got actually a very convincing logic that if that test satisfied, prices will tend to go up, but they might not go up by much. So, is that too harsh or is that too quick to reach a presumption?

So, our answer to that is, no, we've thought about that. And I may or may not convince you on this point, but here's our response, which is, we've spotted the firms with default level efficiencies. I mean, so, we're already building in the notion that the loss of competition is significant in the sense it's more than we would credit with efficiencies.

And if you take the strict consumer welfare standard that's in the guidelines, I mean, the guidelines say, for example, the snip -- the amount of the snip -- I have my quote here -- is not a tolerance level for a price increase. Section -- oh, I don't know -- 1.01 or something.

So, if you take a strict consumer welfare standard, you can say, look, I don't need to know how much prices are going up. If they're going up, consumers are being hurt. And that's what the agencies at least have articulated. But, of course, the Clayton Act talks about a significant loss of competition. So, here that means significant to outweigh the efficiencies we're willing to credit firms with when they propose to merge.

Now, you might say, well, okay, but why not do better? Why not estimate the price increase so you have a sense of whether it's big or little? So, our answer to that is, look, that's actually a much more complicated problem. It requires information on rate at which costs are passed through. The standard thing to look at for economists, pass-through rates, there's a theorem in the paper, a proposition, that shows that you may not have known that you needed to realize the pass-through rate, but that is what you need to do and you need to know in order to figure out price effects.

That’s also true, by the way, for market definition. And that pass-through rate depends on oligopoly behavior and it also depends on the curvature of demand. This is a formula -- a 25-year-old formula from Bulow and Pfleiderer. The pass-through rate depends on the elasticity and the rate in which elasticity changes with price. So, it depends on the curvature of demand. That’s going to be hard to estimate.

We’re prepared to maybe -- we might be able to estimate the elasticity or infer it from the margins, but knowing the shape of demand is a lot harder. Our approach is basically a first-order approach, based on what are essentially slopes and levels around the pre-merger equilibrium. To know pass-through, you have to know second derivatives.

And you can see that here’s another formula for pass-through expressed in terms of the margin and the elasticity of the slope of the demand curve. Again, a second derivative concept. So, that’s going to be very hard to do and we think that’s not practical for a simple test that would be used for the -- to establish presumptions.

I’m going to pick up the pace a little bit, but get through the rest of my slides. So, I hope I’ve convinced you this test is well-rooted in economics --

the underlying principle here is very general. That if costs go up, prices will go up.

We're also not trying to predict the entire model of what determines pre-merger prices. This is one of the difficulties with merger simulation. You try to come up with a model in terms of all the pre-merger prices and then change the parameters. We're just focusing on the change, which is what we care about. All sorts of things go into the level. I don't know about that stuff. It's very complicated. But the change we actually can hone in on and we do that. And we're not drawing boundaries or setting up these algorithms such as one sees in market definition.

The test actually is very general with respect to the shape of the demand system. And those of you who are familiar with Greg Werden's 1996 paper, there's a very close relationship here. I'll glide over that in a moment.

Market definition and merger simulation both depend on the demand shape. We don't need that at all here. That's the big advantage of not caring about -- I shouldn't say not caring, not trying to estimate magnitudes, but only directions.

We also don't actually need static Bertrand. The paper explains this, and there's another paper Joe and I wrote earlier this year about critical loss. If the behavior is non-Bertrand, what you do is you figure out the diversion ratio, the real-world diversion ratio, we call the residual diversion ratio, which is against the residual demand curve.

If I move my price, you may accommodate -- suppose I raise my price. You're going to accommodate your price, let's say. That's going to mean a lower diversion ratio. Suppose you accommodate or a small group of people accommodate, but some other people further afield don't change their prices. Then less of my diversion will be to you because we're moving prices and more from outside.

So, the way to -- so, you can build that in easily by correctly defining the diversion ratio based on the real-world behavior. So, it's actually very general with respect to oligopoly behavior as well as demand shape if you recognize the diversion ratio needs to be what I call the real-world diversion ratio, which is if I'm contemplating a price move and the responses that are really going to happen, what sort of shifting around of sales will occur? And that's something for you to do in the real world.

That's the other thing. We think this is actually easier to implement than what's being done now.

You've got to measure price and marginal cost. This is already done in mergers because you're figuring out margins, you do it for critical loss. I know econometricians tend to think, oh, well, the way -- I'll estimate the elasticity and then I'll infer the marginal cost based on the elasticity. In these investigations, you go and you measure the marginal costs. You've got all this discovery, you look at the cost structure and you can get a good estimate, in my experience, most of the time.

So, prices and marginal costs, and all we're asking to do is measure this stuff at pre-merger levels. We're not talking about predicting some new equilibrium.

Diversion ratio certainly is the harder thing to estimate. It's a key parameter. And, you know, to quote a famous scientist, we want to make things as simple as possible, but not simpler -- but no simpler, and that's what we're doing here. This is the number you need to know about.

And thinking about hypothetical monopolists and -- that's all distracting. This is what you really should care about, okay? And, again, in my experience, this is something you might very well have good evidence on. I discuss more in the paper than I can do in the slides. But whether it's survey data, a firm's actual business documents about who they're gaining and losing, win-loss reports, you know, those sort of things really come up without getting into some artificial exercise of, if all these prices went up, what would you do or what would happen?

When you ask customers these things, they often scratch their heads and say, that's kind of a strange question. But if you ask them if this guy raised his price, what would you turn to? They're like, oh, well, that happened or I've thought about that, here's my second choice. You're much closer to the real world. So, it makes easier to get accurate information.

There's not many things that need to be measured here. And you don't have all this stuff about should you measure things in units or dollars or -- you know, all this peculiar stuff that comes up with market definition and shares that's kind of artificial. None of that.

To the extent you're relying on actual normal course of business documents about margins and diversion, this reduces the scope for game play and litigation. Because you're not asking a new set of questions only in this context. You would look at how the company's actually running their business which, of course, is the preferred type of evidence in any antitrust case. But you're really going for that here directly.

We think it's transparent. There has to be some work done to explain this logic to judges and to the business community. But it's a lot more straightforward than a hypothetical monopolist logic and critical loss. I mean, people are used to that. But if you start with first principles, this is just much more direct.

Capture the notion of loss of competition. We've simplified it so that it's one test formula. Then there's sensitivity analysis. You could say, look, I'm really not sure about the diversionary shift. Maybe it's 20 percent, maybe it's 35 percent. If it's not going to matter for your test, then you don't have to sweat it out so much. If it does matter, well, then it's a closer call and you may need to dig deeper or you may not get a stronger presumption if you're near the edge of the test zone. But you can do that.

And there's no black box here. This is, again, transparent. This is something that would be -- there would obviously be things for experts to argue about in front of a judge or at the agency, but it would be a lot less arcane, I think.

I'm going to skip over these slides that explain the refined version of the test. There is rebuttal then. There's a whole range of rebuttal where the companies, who think about it in court, would say, you know, you measured the margin wrong, you got the diversion ratio wrong. There's some other factors going on in your test. And that would be called direct rebuttal. That's, of course, defined. And then if the companies can't rebut on that, then you move to the full analysis of competitive effects in the back end just as we do in the current guidelines.

How much weight the presumption carries, you know, I think it would have to depend on the test score, just as it does for Herfindahls. So, sometimes that's a sliding factor. But the back-end analysis would look pretty much the same.

If the government or the parties wanted to say, look, in fact, the price effects will be small, we have natural experiments, a la Staples, or we have merger simulation, they can do that the same way they can do it now.

So, sounds good, I think. I'm convinced. So, could it happen? Let me close on this thought. I've gotten some very different sort of reactions to this in the way that I find sort of intriguing. There's one set of reactions which is -- so, from people at the agencies or who do work for merging parties and say -- well, particularly let's say the agency folks say, well, this is kind of what we do, right? We look for these direct effects, you know, for unilateral effects cases. We do this and that's how we figure out what's going on. And then if we think there's a problem, we'll figure out a way to define the market and do all that stuff we have to do to underline the guidelines because that's what the guidelines say and the Courts want markets. So, it's pretty much what we already do.

So, my reaction to that is, that's great, maybe you should put it in your guidelines so they are accurate in terms of what you do, and that would be helpful and we can cut through some of this stuff. So, I'm delighted to hear it. How about some transparency in terms of a process. That's one set of reactions.

The other set of reactions is quite different, which is this is radical, it would be crazy, it would be -- Brown Shoe won't allow it and it would be undisciplined and scary and way too interventionist, whatever. And then I say, well, okay, we can have that debate, but my question is, well, what don't you agree with about it substantively. Isn't it simpler? Is it more accurate? Does it reflect unilateral effects? Is it practical? I think there's certain vested interests in the status quo, in fact, and just inertia, normal inertia, we've been doing it this way for a while.

In many ways, I'm a big supporter of the guidelines. Well, I think in '92 or I'll even go back to '82, you know, big steps forward, big steps forward. But we have had 15 plus years of experience with these now and there's a problem in this area. I think a lot of people recognize it. So, not to throw them out, but to amend them and -- so, I'm hoping this -- we're hoping this will trigger that debate in a substantive way, but not in the sense of, oh, it's a big change, so we couldn't do it and sort of stop there.

And that goes back to the first set of responses. Is it really a big change at the agencies? I'd be curious what people think. It would certainly -- you know, it would be a change for the courts. They'd be getting a different message. And hearing what I heard this morning about Brown Shoe and so forth, you know, I don't think it's necessary to -- one doesn't have to skip market definition.

Take Whole Foods. I think if you would do this, you'd say, look, the market -- we don't really care, basically. The reality is it's either a really concentrated merger in a narrow set of firms, a premium national organic supermarket, or not very concentrated in a broader set of firms. But in that case, it's two firms who happen to offer very head-to-head -- products that are quite close to each other within that broader market because of the product characteristics. So, we don't really care.

If you want -- if the merging parties want and the court wants to pick the broader market, go ahead, pick the broader market because we're not trying to get a presumption based on the shares in that market. So, we'll look at that market and we'll agree that looking at the dynamics and the trends of what's going on in that market is a good thing to do, particularly as it relates to diversion between the two firms, between the two of them.

So, don't get hung up on the market definition. Define a market, get over it, and then go forward. And then, true, you'd have to explain how can it be that it's causing problems even when it's only -- I don't know what it would be in that market. I don't know the numbers. Five percent plus 3 percent or something. I don't know if I'm close. You say, well, that happens when you've got these circumstances. And for better or worse, I think Brown Shoe helped put that point, although I don't really quite go there.

So, that all seems like it could be practical. But, of course, it would require revising the guidelines because the agencies couldn't very well go and do that in court when it wasn't indicated in their own guidelines, both for court and for alerting the business community properly to what would be the policy -- the enforcement policy.

Okay, that's the word.

MR. SCHMIDT: Can we take a couple questions? I don't know. Do we have questions?

MR. SHAPIRO: Only ask a question if you have any doubts about what I said?

**(Laughter.)**

MR. SHAPIRO: Your silence will be taken for assent at this point.

MR. VITA: This may be more of a comment than a question. But to get the process going, don't you need to prove that the guidelines are broken? I mean, I guess that's -- you said we've got 15 years of experience and we sort of know they're broken. I haven't been playing in this game every day for the last 15 years, just watching it from pretty close up, but not playing, have we got a pretty good story for why what we're doing now is broken?

MR. SHAPIRO: Broken's a strong word. So, I hesitate to use it. I'd say there are significant problems in this area. So, I think we can point to certain court cases and we do that in the paper. We talk a little more about Oracle and Whole Foods and so forth. There's going to be a panel at the ABA spring meeting, actually, on possibly revising the guidelines.

I get there by looking at specific examples of where I think it's most clear that the problems are arising. I'm not sure exactly how else we would prove it. I think there's general grumbling I hear around, but that's a proof.

AUDIENCE MEMBER: One sense in --

MR. SHAPIRO: Let me just say one other thing. I think part of the problem is a lot of what goes on never sees the light of day because it's HSR confidential. So, it's very hard -- I mean, 99 percent of stuff doesn't go to litigation. So, you can look at the litigated cases and that's a strong signal, I think, of the problem. But most of it, you can't point to because it's confidential. So, that's why I don't know what else to do other than rely on the experience of the people who do that. It's hard to then systematically assemble that.

AUDIENCE MEMBER: One sense in which I think it is currently broken and it would be a really good thing if we got rid of it is, you know, we have this result that if we predict a price increase enough that we're worried about it, we're implicitly saying that the market is the merging firms, but then the lawyers are all like, you can't go into court and say the market is the merging firms because they'll laugh at you. So, then you have to make up some other market and then you end up with this sort of built-in, one hand tied behind your back before you even start. So, it does seem like there's a fundamental built-in problem.

Aside from the general nature of the whole exercise is kind of silly, there's a sort of litigation problem built into it that I think this would solve.

MR. SCHMIDT: I agree with that and I also think there's the related problem that if you -- take Whole Foods. I mean, it's sort of common sense that, of course, they compete -- I mean, Dan Wells said it really well. He actually represented Oracle in the Oracle case here in February at one of the workshops. He said, you guys are in trouble. You go into court and you say, oh, this stuff's outside the market. He goes in front of the judge and says, look, here's 12 documents of customers who considered, who switched to this. There's some people in competition across these boundaries.

So, if the Court gets in their mind, oh, if it's outside the market, you know, you shouldn't be competing, you're in trouble. So, there's that.

And in Judge Walker in Oracle, he totally got hung up on this point about, well, if you say these guys are going to raise price, then that should be a market and just got kind of tied up in that. So, I think these are very practical problems that are inherent in the current structure of how it’s done.

MR. SCHMIDT: Okay, thanks, Carl. We’re a little bit behind time now. So, let’s take like five minutes to give people a chance to throw stuff away and stretch their legs and come back in five minutes to start the next panel.

**(Keynote address concluded.)**

**PAPER SESSION TWO: BEHAVIORAL AND EXPERIMENTAL ECONOMICS**

MR. LIST: So, good afternoon. I'd like to introduce Paper Session Two, which is about Behavioral and Experimental Economics.

And we're going to do something different here. There isn't a typo in the program. Rob Letzler will be discussing all four papers. And we're doing that on purpose because we added one extra paper to this session. When we put out the call, we received a lot of paper submissions for this particular session and there were a lot of very good papers. So, we decided to accept four rather than three. But as a compromise, Rob will be commenting at the end of all four papers about his thoughts on each of the papers, in particular.

So, after each talk, I will ask you if you have any pressing questions, and if you don't, we will move on and then take all questions at the end.

So, I'd like to start with introducing Stephan Meier, who I view as one of the top experimental and behavioral economists, not only in the U.S. but in the world. So, thanks a lot, Stephan, for joining us and you're up.

MR. MEIER: Thank you very much for the nice introduction and it didn't help reducing my being nervous. But anyway...

**(Laughter.)**

MR. MEIER: So, I'm happy to present this paper, which is, as you see from the title, it's about present-biased preferences and credit card borrowing. It's jointly with my former research assistant, Charles Sprenger, who is now conducting a Ph.D. at UC San Diego.

So, just to start, I'll give you two facts which are going to be important about credit card borrowing. First, it's important in size. So, the U.S. population borrows a lot. If you just look at self-reported data from the Survey of Consumer Finance, they borrow about $30,000 on the average in non-mortgage debt, about 20 percent of that on credit cards, and that's going to be for sure a lower bound because we know that people normally under-report their debt by a factor of two or three, or they outright lie about what kind of debt they actually have. So, it's going to be important for our project because we're not going to look at self-reported data.

The second fact is there is large heterogeneity between people in borrowing on their credit cards. So, again, if you look at the Survey of Consumer Finance, only about 60 percent of people who have a credit card actually carry a balance on it. And our paper is we wanted to explain some of this heterogeneity by focusing in particular on time preferences, on impatient and present-bias preferences.

Now, what do I mean with impatient and present-bias preferences because I use them very loosely? So, what I basically mean is with impatient, that what we normally write in our model, that's the exponential discounting factor. So, that's how much people care about the future.

Now, the second one, the present bias, is there an extra weight on the present when faced with instantaneous gratification? And if you think about the quasi-hyperbolic model written down by, for example, David Laibson, you can think about the discounting factor as the delta and this (inaudible) factor into present bias as the beta which weights, basically, the whole discounting function a little bit downwards at the beginning.

So, what are the effects of present bias theoretically? I mean, first, they value the present so much, it will lead to dynamic inconsistency. So, people might make a plan for how they discount in the future, but when the future becomes the present, they violate that plan and become suddenly more impatient. That might lead to over-borrowing, given their long-run plans. So, they want to borrow actually less, so that's what they plan, but then they borrow too much.

There is a bunch of evidence from laboratory studies that, in fact, people do discount extra -- have like this extra weight or have like present bias preferences. There is a survey in the Journal of Economic Literature that shows that, well, there is a large fraction of people who discount with this present bias parameter.

There are also some new economic studies which can tell you a little bit of a story of where that might come from. So, people might have like two systems, decision-making systems in the brain. One is known as the deliberate system. So, you plan, you think about the future. And the other one is more of an effective system which gets triggered when there is this instantaneous benefit.

Laibson and coauthors show that, well, if you put people in a scanner and confront them with choices very similar to what we confront them with is, if there is this instantaneous benefit, this more effective system actually gets (inaudible).

So, obviously, this has an important implication for IO, how competition works and for public policy, how we think about how to regulate it. However, and this is the prime example -- one of the prime examples in behavioral economics, present bias leads to more credit card borrowing.

Now, the evidence is actually not so great on that issue. And here are two kind of basic empirical approaches so far. One is you take aggregate data and you try to match the moments in that data with each an exponential or a quasi-hyperbolic function. So, for example, Laibson looked at, well, how can you explain credit card borrowing on one hand and holding of liquid assets on the other one, and you're doing a pretty bad job exponentially. You'll do a little bit better if you fit like a quasi-hyperbolic where it has like this present bias parameter (inaudible).

This is great evidence on the aggregate. However, you want to see on the individual level whether those people who are present bias actually have more borrowing. So, that's why experimental economists measure (inaudible) references directly and report it so far and correlated it so far as to self-reported measures of spending patterns, for example.

So, Harrison did, in Denmark, a study where he mainly cares about this long run discount effect. So, he doesn't really look at present bias. He looks at long run discount effects and sees whether people report that they have debt on their credit card. It doesn't correlate in their study. So, long run discount factor in Denmark seems not to be correlated with whether they have any credit card borrowing, self reported.

The second one is a paper in Germany where they do very similar measures of time preferences as we do. So, they are able to distinguish who is actually present biased or not and find out that those who are present biased claim that they have more problems with spending.

Now, this is great evidence, too. But as I said at the outset, self-reported data on credit is a little bit problematic here, and probably more so than when we think about other hypothetical questions because we know they under-report and they misreport substantially.

So, what we do is we basically combine those two approaches in the sense that we measure time preferences experimentally, directly on an individual level. Then we match that measure to objective borrowing data from individuals' credit reports, and then we also control for income and stuff from people's tax data. And I'm going to explain where we got that data from.

So, what we find is, as most other experimental studies, there's substantial heterogeneity between time preferences and who actually is present biased. Long run discount factor, as in the Harrison paper, does not match in explaining who actually borrows and who doesn't. However, present bias is associated with debt problems, and people who are present biased borrow substantially more. This is particular strong for those who actually have a credit card. And Dean probably talks more about commitment devices. That might be one indication, well, there might be some present biased guys around who figured out, well, not having a credit card is actually good for me. So, we can distinguish a little bit there.

So, I'm going to talk about the setup, the results and then I'll conclude. So, what is the setup? We do this study in what is called a voluntary income tax assistance site. So, this is volunteers help earned income tax credit recipients fill out their taxes. So, they come into those tax sites. There are about 22 in the Boston area, and it's run by the City of Boston and the Federal Reserve Bank together, and they come in and they get offered a credit report. And volunteers actually help them a little bit understand what is in the credit report.

We independently measure individuals' time preferences with choice experiment, and I'm going to explain in a second what we do. And then we match this credit data with time preferences, and because it's in a tax site, we also have their tax data. So, we also match it to their tax data.

This was done in two neighborhoods in Boston, Dorchester and Roxbury. In two years, we got about 600 individuals, and for about 540, we had usable matches of time preference.

Obviously, this is not a representative sample, in various respects. One is, they're low to moderate income people. They earn about, on average, $18,000 after tax per year. So, they're extremely, extremely poor. You have to take that into account when we interpret the sizes of the effects.

I think it's more a feature than a problem because we care a lot about people of low to moderate income, because if they make mistakes, it has catastrophic consequences, while if I do -- and I do a lot of them -- it doesn't matter that much. And there is also this growing market for marginal or subprime borrowers, and we are interested in what happens there.

There is an additional selection effect and that is -- so, remember, they come in and we offer them a credit score. Not everybody takes the score and we only observe those who actually have scores. So, we also -- we measure time preference for everybody and see who's selected to that program or who is in our sample. And you see that they're actually more patient, they're more sophisticated guys. So, we should also keep that in mind when we try to generalize from our results to the general population of low to moderate income people or to the general population.

The data comes from one of the credit bureaus. As I said, we get their report. The most important information we use on the report is the amount of revolving accounts. Those are mainly credit cards. Now, if you're familiar with credit reports, that's not debt, per se. That could be convenience charges. We don't know how much of that is actually debt. What we do -- because the Survey of Consumer Finance asks the question, at the end of the month, what do you normally do? Do you pay off the whole amount of your balance or just a fraction and so forth? So, we can look at whether those people who say they pay off the full amount, whether they have actually much more amount on their revolving accounts and they actually have.

You can do the same analysis as I showed you here for the question on who pays the full amount and the results are the same. But we're going to look at the amount for lower income.

There's also information on credit constraints. I mean, we know the limit and we know how much they used of that limit. So, we can control whether our measure of

-- in the choice experiment has actually anything to do with the credit constraint (inaudible).

So, here is what we do. We ask them a bunch of questions. Do you want to have a smaller amount now or a little bit of a bigger amount in the future? So, they go through that list. So, where we ask them $75 today versus $80 in a month, what do you want? Then we ask them the second question and so on.

We ask them a bunch of those questions. Importantly, we ask it in the three different time sets. So, we ask them, today, one month, that was the example I just showed you. Then we extend the period. We say, today, six months, and then we shift the whole period into the future and say, okay, six and seven months. And we use that to estimate or to measure the structure of their time preferences. So, a typical present-biased guy would be very patient in the sixth and seventh month, but if the present was involved, he gets very, very impatient.

Now, you might say, well, that’s just because in the present he actually gets it right now and in the other one, he gets it by mail or whatever and there is some uncertainty involved. Well, we tried to get rid of that. Actually, in both cases, he or she gets the payment by mail. So, we mail it either today and he or she gets it tomorrow, or we mail it in a month or in six or in seven months. Just to keep transaction costs between those two things very similar. We can also look whether they expect to move and it doesn't matter.

So, we pay -- about 10 percent of the participants get paid. And then we can measure this individual discount factor, which is what I called before impatient, and we can see whether people are dynamically inconsistent in those two choices and we use -- in the baseline, we use just a dummy, whether they are present-biased or not. You can also fit a data delta function through the choices. Even though it doesn't fit exactly, the pattern in the data, you can do it and the results are the same.

Now, you might say, well, well, well, what they tell you in those choice sets might have a lot to do with like their credit constraint, which is going to be a problem here. Now, first, others have shown that using those payments to measure time preferences is actually highly correlated to either using primary rewards. So, instead of giving them a little bit less money now or a little bit more tomorrow, you can give them a little bit less chocolate now and a little bit more chocolate tomorrow or choose (inaudible) they are heavily correlated.

You can also look at response rates. So, if you go through that list, there are like choices which are simple. But the closer you get to that indifferent point, the harder it gets to answer the question. So, you can measure how long individuals take to make those choices and that's highly correlated with the measure you get when just using the one (inaudible).

AUDIENCE MEMBER: I have a clarifying question.

MR. MEIER: Yes?

AUDIENCE MEMBER: (Off microphone) Are you doing this by computer and that's how you can measure how long it (inaudible)?

MR. MEIER: Well, they do, we don't. I mean, that's another paper.

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MR. MEIER: Yeah, so what we do -- so, these are other papers. What we do is we can look does present bias correlate with their credit limit? On their report, it doesn't. For a subset of people for the 2006 sample, we also get their credit report one year after the experiment. So, we measure their preferences today and then see whether we can predict how much revolving balance they have one year later, just to get a little bit rid of that mighty shock (inaudible) and the results hold. And you can just -- in all the regression, you can control for a limit and their FICO score and -- one or the other, and the results are not affected.

So, what are the results? So, what you see here is outstanding balances -- these are the raw correlations. Outstanding balances whether they're present biased or not, and you see that there's a substantial difference of about $700 in what present -- and, remember, they have $18,000 in disposable income per year. So, we think that's a substantial difference in what present biased guys carry on their credit cards and people who are not.

Now, you can control -- if you basically look at column two, where we control for some social demographics, in particular income, the number of dependents. That's from the tax information, their educational level and some demographics. And you see, first, that the individual discount factor, that's basically the exponential discount factor we normally write down in our models. It doesn't do anything to their borrowing. Very similar to what Harrison found in his study in Denmark.

However, if you look at the present bias, present-biased individuals carry higher debt on their revolving accounts. Those are two (inaudible) so it's hard to interpret. If you knew the marginal effects here the probability that you have any debt increases by about 14 percentage points and then conditional on having debt about $500 more.

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MR. MEIER: No, no.

Okay. Now, you might worry what I said before. Those are -- here, we also have people who don't have any credit cards. But you might say, well, this is actually choice. That's why we don't exclude them in the basic regression. But you can also say, well, let's only take those who have credit cards and control for the limit and we can also -- because we don't have interest rates on their debt, which would be an interesting variable to have, obviously, we can control for that FICO score to get a proxy for what their interest rate might be if you believe in like risk-based pricing in those markets.

Now, the result, again, look at column two, becomes stronger in the sense that if you get -- if you look at marginal effects here, the probability increases by about 20 percentage points that you have any debt and by about $1,000 conditional on having debt.

Now, we can do a couple of robustness tests. You can use alternative measures of dynamic inconsistency based on those choices. You can also -- data (inaudible) is one of them. You can include risk attitudes because you might be worried -- well, you might pick up some of the curvature of the utility function in those choices and not just time preferences. It could be a question on the risk attitude and you can control for that. You can control for whether they expect to move, because he might say, well, you ship it, but it might be a problem if they expect to move in six months, that the probability that they actually get it is very different between the different choices.

We can get (inaudible) and include multiple (inaudible). I hadn't talked about that. For some people, that's very hard to calculate time preferences, because what they basically do, they switch around. So, they say like, well, 75 over 80, well, I'll take 75, and then what about 70 over 80. I'll take 80. And 65 over 80, well, 65. So, they seem to -- I don't know exactly what they do. You can take their first switching point and assume that's really the one they wanted to and the results are the same.

Again, as I said, it's very important for one sample, we looked at their borrowing one year later, the results are the same.

So, let me conclude. I think we tried to combine experiments in the sense that the methodology from lab experiments to measure people's time preferences and combine it with behavior outside the lab and the results seem to indicate that the structure of time preference seem to explain neither outside the lab, the long run discount factor doesn't do much. However, present-biased preferences seem to matter.

Obviously, this has implications if we think about competition, and there are a couple of models who have already put present-biased consumers into those models and they seem to show that competition doesn't work as well as you get when you assume that people are exponentially, and it might have implications about how to protect if you want those guys who seem to make a mistake in the sense that they want to be patient in the future, but they aren't actually when the future becomes the present.

I'm doing some more work on present bias and credit contracts in general. So, a very similar study we're conducting right now is to figure out whether we can explain why people pick up adjustable rates over fixed rate mortgages. And so, we do a very similar analysis to that in the mortgage market. And I didn't talk much about (inaudible) and commitment mainly also because Dean will talk much more about that. But that's, I think, the next step to go.

Well, how many people of those who are actually present biased know that they're present biased and take actions to limit their vulnerability, so to speak, from them? So, thank you very much.

**(Applause.)**

MR. LIST: Okay, very good, thank you, Stephan. Are there any pressing questions at this point or can we move on?

**(No response.)**

MR. LIST: Okay, great. James, why don't you come up and load. I'll give just a brief introduction.

So, our next paper is by James Hilger. Obviously, I just met James today. And he is talking about an experiment that he's run in one of my favorite markets, the retail wine market. So, he'll talk for a bit, and if there are any questions, I'll take those at the end.

MR. HILGER: Thank you. Thank you, everyone. First, I have two housekeeping things I need to take care of. One is that I have the typical FTC disclaimer. These are the views of my own and don't represent the Commission or the Bureau. And the second is that if you have a copy of this paper, today, I'm going to present recent results. They are qualitatively the same. So, if you get lost, if you're flipping through the paper, that's why.

So, I'm going to talk, as John said, about expert opinion labeling in the wine market. So, experience goods are goods that are defined as products or services that you don't know the quality of that product or service until you've actually consumed it. And in today's marketplace, there are a lot of goods, in fact, one might say most goods, that you really don't have a sense for what you got until after you've had it. Wine is an example.

Books could be an example. A dinner could be an example. And you could stretch that to some of the things I work at here, major appliances or lightbulbs. One might not know the impact of a purchase on their electricity bill until after they've bought that product and experienced the good.

So, there are a lot of areas where consumers rely on the opinion and information provided by experts to make their decisions, and I'm going to look at the wine market, specifically.

First, a little background, Jen and Leslie have a paper that looks at the impact of restaurant sanitary quality postings. And they found that consumers respond to the higher-quality, A-grade restaurants. But this paper doesn't really get at some of the aspects of what I want to talk about. It's hard to separate out the information -- the actual provision of information, the quality that’s posted and the actual quality of the restaurant.

There are a lot of different sources of information that consumers might be using to decide if a restaurant is of high quality or low quality besides just the letter score that’s posted on the door. For example, if it’s crowded and, you know, just the general look of it possibly. So, there are a lot of sources and I aim to, basically using the experimental approach, account for this spurious relationship between price and quality.

Another paper by Reinstein and Snyder looks at movie reviews. And they find, looking at Siskel and Ebert reviews, that there’s no significant impact of movie reviews -- positive movie reviews on revenues in general. For some subset dramas, they found that there’s a positive effect. But, in general, you know, a movie review, it wasn’t -- they couldn’t find a significant impact that increased revenues of the movie. Sorenson and Rasmussen look at book reviews and they find any publicity is good publicity.

So, using the experimental approach, I’m going to basically try to separate out the impact of both the expert information being provided, which I’ll call the promotional effect, and the actual information in the expert opinion, which would be the information effect. And that brings me to the retail supermarket where consumers are often -- they have to decide, in the case of our stores, over 700 UPC codes, 700 different bottles of wine, and she really wishes that she had some information in the store. So, that's a photograph of the mock-up label that my coauthors and I, Greg Rayford and Sophia Viabose (phonetic) put in two Safeway stores in California in the spring of 2006.

So, we're going to investigate using a difference and difference methodology, whether or not there was a statistically significant impact on demand of the label and whether or not information -- there's evidence that the quality signal was used.

So, data, I have data for over two years, scanner data for 41 stores in California. It says detailed wine information, price, whether or not it was discounted and the sales data. Then we also compiled a database of expert opinion scores from the major wine reviewers and we drew a random sample of wines that were in our treatment stores to put up in the treatment stores from the set of wines that expert opinion scores were available. So, we have a random sample here, but it's not a random sample of the wines that were in the store. We have a random sample of the wines in the store for which scores were available. So, I just want to point that out.

As I mentioned, we have a lot of different control stores. This is a unique problem that in the dif-dif literature, you know, we had more stores than we wanted to use. The stores were all in the same geographic area, but, you know, they had differences in sales. So, we wouldn’t necessarily want to pool all of the stores together because the consumers in different stores might be systematically different.

So, to select our control store, we went through several different analyses. One just looked at the demographics of the consumers in the area, of the store. We also ran estimated demand -- reduced form demand equations and then did Chow tests for pooling, and we also looked at if the store sales across stores moved in parallel over time.

After all of that analysis, we constructed a control store set of 13 stores, which most all of these tests indicated that they were comparable to the treatment store. And we did this twice because we have two treatment stores. But I’ll only be focusing on one of those stores.

So, the empirical strategy is basic dif-dif. We also did dif-dif indifferences, and for the sake of time, I'm going to skip through that. I do want to note that in this slide K would be whether or not the actual wine was labeled. So, we have T is a dummy for the treatment time period; J being the treatment store; K being a wine that was actually labeled. And this dif-dif analysis we did over two time periods, not the two years, but two months.

AUDIENCE MEMBER: (Off microphone) Can you tell us why you'd need a regression if you have an experiment?

MR. HILGER: Good question. The question was, why do I need a regression if I have an experiment? Because there are some covariates and when I include those -- so, the first model is just basic dif-dif and dif. And when I move on, we'll find the impact of the heterogeneity in one of the covariates.

AUDIENCE MEMBER: (Off microphone) (Inaudible) experiment or (inaudible)?

MR. HILGER: It's a field experiment.

AUDIENCE MEMBER: (Off microphone) Wouldn't you randomly assign (inaudible) information on pricing conditions across bottles of wine?

MR. HILGER: Correct.

AUDIENCE MEMBER: (Off microphone) So, why don't you just do it with (inaudible) across different bottles of wine and you don't have to worry about doing (inaudible)? I guess that's what I'm -- it might be useful to clear that up for the audience.

MR. HILGER: Okay. The question is, why don't we just compare -- you know, put the expert opinion label up and then look at the impact in the one store between wines that were sold and wines that -- I mean, wines that were labeled and wines that weren't labeled. Well, one thing is if we did that analysis, we wouldn't be -- we might find a shifting of the -- you know, if people -- you might find a switching effect.

Also, we wanted to control for previous time periods and it's not clear to me, you know, at the moment quickly how those -- the time switches and time trends and trends across store might impact that. But, most of all, we wouldn't be able to -- well, let's move on and maybe address that.

So, the first results, on the left is a triple difference and on the right is the dif-dif. So, in the top red highlighted box, we have a store month effect, which is basically did the treatment store sell more wines in the treatment period compared to the control store set, difference between treatment and control periods. And we found that that was a positive and significant effect, which was, you know, to be clear, somewhat worrisome because this is the effect, including the effect on wines that weren't treated.

So, the treatment store and the treatment period saw a relative increase in sales, even on the untreated wines.

We have a positive effect insignificant on label, store and month, which are the treated wines. I do want to note that less than 1 percent of the K dummy -- of the L -- I mean, this label, store, month variable, less than 1 percent are one. So, you have a serious power test and the probability of getting a T statistic that significant is extremely low. But there is a positive effect.

Then when we look at just the labeled wines, we keep the positive effect on store, month. So, this is just the labeled wines, but it's not significant. So, the upshot is the average treatment effect on treated wines is positive, but not significant.

Keith?

KEITH: (Off microphone) When you say labeled wines, it's not information on the label, it's the one you put a (inaudible) on?

MR. HILGER: Right. So, in this, we are not controlling for the information yet, this actual score. This is just it received a score.

Now, in the next model, we're going to include some of the wine covariates, such as the actual score, the price, and I'm going to note here that price is always negative, because I have an interaction term and it was easier to deal with that way, in a promo, whether the wine was on sale, and a dummy variable for red wines.

So, I've run this several different ways, building up from the most basic model and in several different functional forms, quantity logged and not logged and price and the score logged and not logged. What we find when we include the score store month -- so, this is the impact of -- the marginal impact of the score on a treated wine in the treated store during the treatment period is positive and significant for all of those models.

Well, I should state that the average score treatment effect is positive once you keep in mind or take into account that the average score on a bottle of wine was 84. So, if you calculate it out, you find that the wines that, on average, you know, the high-priced wines -- well, on average with an 84, there's a small increase. But high-priced -- I mean, high-scoring wines saw a large increase in sales of roughly -- well, depending on which model you look at, you know, roughly eight bottles. And I should have mentioned that the average quantity sold was about nine bottles. So, they saw roughly a -- not quite a doubling in sales, but they saw a fairly large increase in sales.

In wines that received a low score, less than 70, saw a decrease in sales.

AUDIENCE MEMBER: (Off microphone) (Inaudible) zero if they weren't labeled?

MR. HILGER: If they weren't labeled, I didn't -- so, these are only on treated wines. This is only on the subset of treated -- of labeled wines.

AUDIENCE MEMBER: (Off microphone) Is the score variable a continuous variable?

MR. HILGER: Score is continuous, right. Score is continuous. Okay. So, one last thing, we also estimated this model using all untreated wines. I mean, just the same model, but we didn't include the treated wines in the model, which is the second column.

The third column, we did it for a placebo store, you know, just a random store from our control group and used that as a treated store, and also a placebo time period. In all of these, we wouldn't expect the result to be significant and it was not significant.

In conclusion, we feel fairly -- extremely confident that in the two Safeway stores that we had this experiment, there was, you know, a positive, but not significant, marginal effect on having the wine being treated for the average wine, but scores that -- wines that received a high score saw a significant increase in sales and wines that saw a low score saw a negative -- I mean, a negative change in sales.

So, this is -- you know, evidence points to the fact that consumers could actually utilize the information that's posted on the label, which is interesting. You know, there's evidence that they're basing their purchase decisions on information that's provided, which was what we sought to investigate.

I think I'm out of time or over time, so I'll wrap up now.

**(Applause.)**

MR. LIST: So, I'm going to go on because James got a few questions from us that were sort of spontaneous. But we'll come back to that one at the end for anyone who has any other questions.

Our next paper will be by Cary Deck, who is a dynamic, young experimenter from the University of Arizona. Cary and I met at the University of Arizona when I was a faculty member there. Cary will be talking about price discrimination with sequential purchasing.

MR. DECK: Thank you, John. So, first, I guess I should acknowledge that this is joint work with John Aloysius and Amy Farmer, who are both at the University of Arkansas with me.

The other thing I guess I should point out, since everybody has gotten up here from the FTC and made this big disclaimer of not speaking for the FTC, even though our college is named I guess from a certain Bentonville-based retailer and they have given large amounts of money to both the CRE and the ITRI and the college which funded this research, I, of course -- they never talked to me.

**(Laughter.)**

MR. DECK: It would be nice, but they've never talked to me, frankly.

So, this paper is a little bit different in that I'm not trying to explain anything that I know is going on currently, although it may be. I just have no information on that. But I'm trying to be a little more kind of forward-looking in thinking about what might be coming down the road.

And, so, if you think about going through any large retailer, the firms are currently using RFID technology, little radio frequency tags, at the pallet level to track kind of large shipments from the wholesaler to the back room, putting them out on the floor, and then when it goes to the crusher.

They're starting to include item level tags, which allows them to see individual items, when they come through the storeroom out to the floor, when they're picked up by a shopper. The main uses currently are for theft detection, so to see if somebody's trying to sneak through the front door with this item in their pocket. They can track to see if that's going on. And also to be able to restock. They can tell sooner that product is off the shelf and they can then restock the shelf more quickly.

But another thing that this would do is allow you to start setting prices on items based on the other items that are in somebody's cart. So, you could think about a seller might know that -- since they hear that you've already got some peanut butter in your buggy, when you turn to go down the jelly aisle. So, they can see, look, this person has brought crunchy peanut butter and they may be able to tell something about what you're likely to want to purchase on that next aisle. What kind of things you're likely to want to consume next.

If you go and pick out a shirt, then they have some idea about your tastes. When you go to pick out a pair of socks or a pair of pants, they would know, well, ones that match this or coordinate with it are more likely to be of higher value to you.

So, if you think about how you might try to set these prices on different goods, within a shopping trip, maybe it's not too hard online. You can track pretty well what people have looked at, maybe where they've searched, what items they've viewed and not put in their shopping cart, what items they've viewed and had put in their electronic shopping cart. In the store, it may be a little more difficult, but there's new technology that's coming along, like RFID, like this smart shopping card, that would allow you to do more of this kind of practice.

So, the idea here is there's this little shopping cart add-on and you can see the price on any item to keep a running total for you to throw it in your basket. You can certainly envision a world where it says, hey, we see you've just bought this, and like Amazon, recommend some other products you might want to buy. You've bought this, maybe you'd like that, too. You've got four of the items for a lasagna, don't forget you probably need to buy this item, too.

So, it's going to be possible to start tracking these purchases as you move around the store and start setting prices to individuals based on the items that are in there. So, you can think about these little things popping up, little coupons. And say, oh, this is great, I've got a 50-cent off coupon when I approach the jelly, that's wonderful. Not knowing, of course, that the next person got $1.50 off the same jar. I'm not going to be mad, I got a coupon. They're not going to be mad, they got a coupon. And we don't know that each other is being treated differently here.

So, what we want to look at is to try to figure out what the implications of this might be. So, sellers now have large amounts of information on buyer preferences. They get scanner data if you have frequent buyer cards. We use credit cards. They can track your purchases across time, everybody else's purchases across time, and they can get a pretty good idea what kind of goods have what kind of relationship with each other, whether or not the values of particular goods are highly correlated or not, whether or not the goods are compliments or substitutes or not. And, of course, those concepts -- sometimes when people think about them, they view them as the same, but they're very different concepts. It can be correlated, but not be complementary or substitutable and the opposite is true as well.

So, traditionally, what sellers have had to do is basically set their price in advance, but now we're thinking about what happens if the seller can adjust those prices in real-time. If we want to think about what monopolists might do in this situation, it's kind of a nice standard starting point for thinking about pricing. And we want to think about what might happen in more competitive markets with this kind of ability.

So, traditionally, in what you could also label as pure components, the sellers maybe know the distribution of buyer types and they set optimal prices. They're going to post prices prior to the buyers coming in and observing and making that purchase decision. So, the seller is going to attempt to increase profits by doing price discrimination like quantity discounts, coupons, all of these types of things. They're kind of more generally applied.

Another technique, at least in thinking about combining different items, is mixed bundling. So, what you do is basically sell the components by themselves and allow the person to buy the bundle of the two goods. So, going back, you know, for a long time we've known that you can generate increased revenue, at least if you're a monopolist by basically jacking up the price on the single item and then cutting people a break on the bundle. So, this is a technique that people use to try to -- sellers use to try to increase those revenues.

There's a recent paper by Venkatesh and Kamakura where they try to basically go back and redo the Adams and Yellen type of setup where they exam explicitly kind of the degree of complementarity, that theta there, the degree of complementarity between the individual items that are being bought. So, there's good A and there's good B. The buyers have values for those single items and then they have the value of the bundle from buying both items that some combination or that some multiple of the sum of the single item prices.

So, we're going to follow along this -- ultimately, we're going to follow along this same kind of model structure here and trying to think about how sequential pricing might influence the market.

So, with sequential pricing, just a little example to drive home what we're saying here. So, the seller actually can set prices incrementally during the shopping trip. And they're setting them to a specific buyer who is there. So, we're thinking about having two types of people and there's a high-value person, Person X, I guess, who's got a value of 100 on both goods and there's a low-value person, who's got a value of 20 on both goods.

If we assume that the bundle value is just additive, there's no complementarity or substitutability in the products, if under fewer components, well, the seller is going to basically set a high price on both goods. The high-value buyer is going to buy both and that's the end of the story. They're going to expect to make $100. They don't know which type of buyers is visiting. They could drop the price to $20 and guarantee selling to everyone, but that's lower -- guaranteeing $20 is lower than getting $100 on average.

With mixed bundling, the result's the same, because there's not these people who have high values of one good and low values of the other good that you can kind of exploit. It ends up being that this mixed bundling would generate the same level of profit.

But if you could price them sequentially, you walk in, you've got a price of $100 on the first good. Buyers say, yeah, I'm going to take that; no, I'm not going to take that. A person who takes it, boom, you know what they're willing to pay on the next item and you can charge them a high price, whereas the person who refuses to buy it, you can then charge them a lower price on that second good.

Now, there's a clear assumption here which is that you're preventing them from knowing the second price until after the first price has been revealed and they've actually made that decision. So, we're assuming that they have, in fact, committed themselves to buy the good once they've taken it, before they see the second good's price.

So, with sequential pricing, it’s still not the condition that you’ve got to have discrimination going in that second market. There may be some information that’s revealed to you that you can just use that you would charge to all buyers in that second market there. But, basically, the seller posts a price for good A; the buyer makes a purchase decision on good A, then the seller is going to set the unconditional price of good B, unconditional in that it’s not based on whether or not the buyer purchased good A, and then the buyer can make the purchase decision for good B.

So, a good, standard monopoly story, you go to the end, the second good market. The seller knows that some people have already bought A, some people have not. They may have different values, depending on the theta that’s there, if these goods are complements. People who have bought the first good are going to have kind of a higher distribution of values for the second good.

And then, based on the answer there, given this optimal second stage price, back up and figure out what’s the optimal price to set in the first stage. Of course, here, what the seller is trying to figure out is, I’m going to set my first period price in such a way that I can -- I get the maximum exploitation at the second stage. So, if I’m worried about the profit on this good, but I'm also taking into account how I'm going to exploit them next time based on what they do here or based upon how their values might change.

AUDIENCE MEMBER: (Off microphone) When you say in stage three unconditional, you mean you're going to set the same price to all consumers, but we're going to condition it on the information that (inaudible)?

MR. DECK: Well, what I mean here -- what I mean here is that they're setting the same price for people who did and did not buy the first good.

AUDIENCE MEMBER: (Off microphone) That's all you mean by --

MR. DECK: That's all I mean by this.

AUDIENCE MEMBER: (Off microphone) (Inaudible) using the information you know (inaudible)?

MR. DECK: You know about the information on the two goods. So, now, you can take into account if I set a low price of good A, lots of people buy it, it may affect their values for the second good. I'm using that information. So, in particular, in the goods or say complements where theta is positive, the people who bought A now are going to have a higher value. So, when I'm setting that second good's price, I may want to push it up if there's a lot of those people that I can exploit.

So, to go any further than that, we need to make some kind of assumption on the distribution of values. None, in particularly, seems to jump out to me empirically as something great from the current markets. But a uniform distribution is a nice convenient trackable way to start. So, this is a general assumption we’re going to make.

So, if we assume the values we just distributed uniformly kind of over that square, we can work through and determine -- you know, it’s a long algebraic exercise, but you can work through and determine what’s the optimal price to set of the second good, and given that, what’s the optimal price to set at the first good. I don’t show it up here because even in this simple problem, it’s -- just the answer is about a page and a half long. So, there is a complicated answer one can write up there.

Let me just point out if there’s no cost, if marginal costs, which is C, are zero and there’s no additivity in the products, then you’ve basically got two unrelated markets and the firm ought to set monopoly prices in both markets. So, that’s that result.

If you can discriminate, now, step three becomes that you’re going to set a conditional price. So, now the monopolist is facing these two concerns. One is what price do I charge people who did buy the first good and what's the price I should charge -- the second problem is what's the price I should charge to anyone who didn't buy the second good. And then given those, take that into account in the first stage, they find the optimal good A price.

Again, under the same uniform distribution assumptions, this works out much nicer. You can figure out what the optimal prices are at each point, for each type of person. And, again, if there's no marginal cost on the good and theta is zero, so the goods are just -- the value of the bundle is just additive, you would charge the standard monopoly price in both markets. I mean, there's no information coming in from whether or not they bought good A or not. That doesn't tell the seller anything.

So, what I did do is just kind of go through and do some numerical comparisons for different thetas to see what the implications -- so, we did some comparisons to see what the effects of theta were. Basically, we compare this with pure components, mixed bundling. It turns out when the goods are substitutes, this kind of practice can be very effective because, basically, it prevents the buyer from substituting the first good for the second good because you're kind of holding back the second good on them.

We looked, also, at the correlation between the goods. I'll say a little bit more about how we did this. Basically, the distributions we used were just removing the corners from that 100 by 100 square. So, if the goods are highly correlated, like the example we showed before, sequential pricing can outperform mixed bundling. So, it can be an effective tool, but it's got to be at a pretty extreme level before it beats something like mixed bundling.

So, now, what we want to do is think about a competitive market --

AUDIENCE MEMBER: (Off microphone) Before you move on, are you assuming that buyers do or don't behave -- are you assuming they don't behave strategically, the buyers (inaudible)?

MR. DECK: In this case, they don't. The buyers want a single unit of each good or, at most, a single unit of each good. They're going to make one choice. In the case of sequential pricing, they don't know what prices are coming.

AUDIENCE MEMBER: (Off microphone) But do they know that the choice they make in the first -- is there just no room for them to behave strategically or you're assuming that they don't know (inaudible)?

MR. DECK: So, one could think about how they would react if they anticipated their decision at the next stage. So, we could think about -- we're abstracting away from that. We're just assuming they're not. But you could think about this as a brand new product, one they didn't even know existed, so they've got no reason to be formulating a price expectation on it.

In the interest of time, since I'm evidently way behind, we're going to look at competition using experiments, going to the laboratories, seeing how firms might behave in this case. We're going to introduce the idea of informed and uninformed shoppers. But, basically, some of the shoppers know the price of good A at every seller and some of them only go and visit one seller when trying to make their purchase decision.

Now, this is actually a really important point here at the bottom that we're assuming only buyers who actually visit seller I observe that person's price. This makes the problem much, much uglier, but it also seems like a much more realistic assumption to us.

I could envision ways that somebody might be able to comparison shop, go to an internet website and say, oh, here's the prices that people are offering for good A, oh, I've read ads in the newspaper, here's the prices that these people are promoting. But if my value is such that I'm not going to buy that good, it's hard for me to understand how I would go to the store anyway to then be shown what the good B price was. So, an inherent part of this is that the sellers are withholding good B information until the good A market is clear.

We use a basic three by two across subject design. There's three kinds of market value conditions. Baseline condition where theta is zero, there's no complements or substitutes there; a complements case where theta is .3; and then a case with correlated goods. And, basically, what we do is just lop off kind of two corners of that unit square. I've got a little graph here. This was an online tool that the subjects had when they were making their prices, studying their prices. So, they could actually see -- they could see for any set of prices they wanted to pick, put up here, we would tell them what a monopolist would earn.

So, if you did this, here are the different people who would buy good B only, good A only, these people would end up buying both. You could click on here and see what the values those buyers had at a particular point were. But to get the positively correlated goods, we just lopped off kind of the top right corner and the bottom left corner. Not pretty mathematically, but it's very nice and very easy to explain to somebody who doesn't understand what a joint distribution is and all of this. We're going to draw from this.

A few other things. We've set marginal costs to zero. We have four sellers in these markets. These sellers are undergraduates who are in the role of a firm and they're paid based upon the profits that their seller earns. These took about 90 minutes. They ended up making about 18 bucks. There was a lot more money at stake for them. But as you'll see in just a second, they were just very, very competitive.

These markets went very, very fast. The buyers are assumed to be non-strategic. One buyer kind of enters the market, reveals their decision and leaves. So, the buyers are all automated. They're just computerized. They show up every three seconds, see the prices. Depending on whether or not they're informed or uninformed make their decision and they leave the market.

The sellers can adjust prices at all times. They can observe perfectly what their rivals are charging and, of course, they know all the parameters as well. They know the distribution of values. They know how the bundle value is created. They know what percentage of people are informed. They know their rival's cost. They have complete information far more so than would occur in a normal market -- in a naturally-occurring market.

So, just quickly kind of the base results here. I'll just show the figures. There's econometrics in the paper if you want all that detail. But, basically, in the baseline case, the ability to discriminate doesn't really seem to influence the prices that the sellers charge. So, whether or not they're allowed to discriminate or they cannot discriminate, it doesn't seem to change what they're doing there.

The fact that they don't end up discriminating based upon whether or not -- when they can, they don't end up discriminating based upon whether or not the buyer bought good A, which when theta is zero, they shouldn't. I mean, there's no information in that and, therefore, you wouldn't expect them to charge those buyers different prices and they don't.

I'll just point out the ability to price discriminate does not affect welfare here. But what you can see is that the good A prices are way down low. I mean, in fact, we have multiple times where people were giving away good A because they were trying to capture the good B market and get those comparison shoppers to come to them. So, they were actually setting very, very low prices on good A. And then, of course, where there's no competition for good B for the sellers -- or for the buyers who have come to them, they're charging much higher prices. Theoretically, they should charge 50 to everyone who comes to them, but they don't.

AUDIENCE MEMBER: (Off microphone) When you started, my reaction was you must (inaudible) buyers, and then when you go through it, it looks like you have (inaudible) sellers. But I realized that you're giving them three seconds to make a decision and I wonder how realistic that assumption is because I would think that business firms planning to use this kind of strategy would have a lot more time to think about it. But then if you have any kind of experience where people make repeated purchases, they can operate strategically.

(Inaudible) shopping with my wife, if I had any indication that firms were making these price discrimination decisions so that you would both compare prices that we would offer, only one make a purchase and then go on to the next good and see if we could get around the discrimination. But you're coming up with price discrimination on (inaudible). I'm not sure that if the businesses had more time you would still get that decision. But, even so, I think smart shoppers (inaudible) so that you know the result.

MR. DECK: Well, let me respond because you went through a lot of things there.

MR. LIST: On the other side, the comment was you must have stupid sellers to do this, but then you said, well, they only have three seconds to make up their minds how they will price.

MR. DECK: Okay. They can adjust their price at any point in time. Buyers show up every three seconds. So, it's not as though, you know, Wal-Mart gets a long time or anybody else gets a long time between when people arrive at the store. I mean, they can adjust at any point. They go for 750 periods. It's about an hour, right? If you think about prices being set daily, this is a couple of years' worth of experience greatly condensed, but they also have a lot of information.

Now, the second part about whether or not they price discriminate, well, in this market, they shouldn't. If we go to a case where we have complements, in this case, we're, again, seeing very low good A prices, but they are price discriminating. So, people who are buying good A and can be identified as such are getting a much higher price than the people who did not buy good A. So, they are discriminating along those lines.

What's interesting is that the people where the discrimination is not possible or you can think about them not being able to identify whether or not the person's bought good A, those buyers are getting the same prices, at least in aggregate, as the people who could be identified as buying good A. So, sellers are, in essence, treating everyone as though they're the high-value type buyer, as though they're the person who bought good A when the sellers can't tell. That's the result that we're getting off of this screen.

We also here find that the ability to price discriminate does not affect welfare. But part of what's going on there is that because the good A prices are so low, almost everyone's buying good A. So, that's a fairly small segment of the market that's not buying good A, which is why the welfare effects aren't very strong there.

With correlated goods, we, in essence, get the same kind of thing. The buyers -- when buyers can be tracked and identified, people who are known to have a high value are getting a much higher price. When sellers cannot identify, they're giving everyone that same higher price.

So, I'll just compare this real quickly, since I'm out of time or over time, to some experiments that we ran just with bundling. In these experiments, people only could price discriminate or not. They were doing sequential pricing. Here, they were bundling or not.

These experiments were far less efficient than with bundling and the sellers made a lot more profit here than they did with bundling. So, just to summarize, this new technology could be very useful in a lot of ways, making recommendations, providing new information.

It appears that the ability to discriminate is going to kind of dampen the effect of doing this kind -- I'm sorry, the competition will dampen the effect of the ability to price discriminate in such markets. But just the ability to kind of price sequentially there may, in fact, have harmful effects. And I will stop and answer questions later.

**(Applause.)**

MR. LIST: Thank you very much, Cary. That was very good.

Our last presenter is Dean Karlan from Yale. Dean's a co-author of mine and one of a group of scholars who is taking field experimental methods to important issues in development economics. So, I think this is an important line of research. So, Dean, fire away.

MR. KARLAN: Thank you, John. So, since I'm last and shortened for time, I'm not going to actually trim any slides. I'm just going to talk really fast.

**(Laughter.)**

MR. KARLAN: So this is joint work with Xavier

Gine and Jonathan Zinman. The heart of what we did is really the perfect segue from Stephan's talk, the first talk, Stephan Meier, about time inconsistency. What's the implication from time inconsistent models whether it's from -- any one of the kind of pick your model from a theoretical perspective or we're talking about models of dual self, hyperbolic models, quasi hyperbolic models, whichever model you have that predicts or makes the statement that there's some inconsistency over time in the way people behave.

They all share a very common prediction, which is that people should have a preference for commitment, that when we look at the world around us, we actually want certain things to cost us more money. And this is not the normal way that we have traditionally thought about a lot of situations.

You might actually -- different people might have different facets of their life in which they have this preference. You might prefer that peanut M&Ms cost $100. I know I do. Other people might just prefer that the cost of debt was radically more expensive so that they wouldn't borrow. But you don't control -- I can't get M&Ms to jack up the price of peanut M&Ms. It's not in their interest to do that. So, how does one go about doing that and are there products that can be offered individuals that effectively have that element to them? That allow them to raise the price of things that they want to be more expensive.

And, so, we did this with smoking. The very simple question is, if there's people who want to stop smoking, but they basically find themselves too tempted in the future, and so, when the future is now, they end up smoking. So, are there people who would say, you know what, I would really like it if you could get cigarettes to be more expensive?

Now, you can't actually go around and raise the price of cigarettes except unless you're going to do it through the government and do taxation. But, you know, what is a private market solution to this is by having people put up a bond, put up a contract that basically commits them to stop smoking. So, I'm going to explain how we go about doing this.

The one other element that Stephan referred to as well that is a necessary element here, it's not just time inconsistency, but there has to be an element of sophistication. I have to not only want higher -- I have to not only want to stop doing something or to start doing something else, but I have to know enough about myself to know that at the current price scheme, the market prices that are going to be out there in the world, that I'm going to screw up. And I have to be sophisticated enough to want to actually change those prices. That's a different element and that's not something that's so obvious how we go about identifying who's whom in that spectrum. Because, clearly, there are going to be people who would be naive about this and think that they will actually change their behavior despite even if prices don't change.

So, I have skipped a whole bunch of stuff that I will talk about. Other than to say I'm going to talk a lot about smoking, because that's what this project is, but the general concept can apply to savings, can apply to borrowing, can apply to exercise, weight loss, voting, and we actually have a Web site that we created in the United States for doing just this, called STEKK (phonetic), with an extra K. The extra K is for contract.

So, what we did in the Philippines is we created this product called CARES.

No, no, no, no, in legal -- that's what I'm told. I remember this from being a child, taking notes from my mother when she was in law school, that contract is written with K in law school. So, the product -- I told you, I'm just going to talk fast here now.

So, CARES is called Commitment Action to Reduce and End Smoking. In the Philippines, acronyms are popular for everything. Everything has to have an acronym to it. Other things we've done are similar in this spirit, that they always have these kind of catchy acronyms to them. The commitment savings account we did was called SEED.

So, you open up an account with 50 pesos, which is $1.25. We're basically dealing with people -- this is not -- you know, it's a relatively poor area of the Philippines, but this is not the poorest of the poor, by any means. It's in the southern area of the Philippines in an island called Mindanao.

A bank went out into the field and offered individuals a bank account. And they said, look, here's how this account works. You put in a dollar to start the account. You have to do that. We will then come to you every week and collect money from you and you're supposed to -- we're going to give you a little box. This box looks like this. This is where you're supposed to put the money that you're putting into cigarettes, instead put it in here. We'll come by once a week. We have a little key for this box. We'll open the box up; we'll empty out the money; we'll take it; we'll deposit it into an account.

This box, to be clear, could easily be broken with any simple sledgehammer type device. It’s not a foolproof system. It’s just a mental account with a small physical barrier.

Now, at the end of six months, the bank officer comes back and has them take a urine test. If they pass the urine test, they get their money back, zero interest. The reason for zero interest is very simple. The bank wouldn’t do it with interest because -- why? Because, otherwise, they’d be giving away a free deposit collection service and they would have a bunch of people who were not smokers taking this product up, and they knew that that was just not the way to run this as a business.

So, if they failed the urine test, what happens? The money goes off to a local orphanage.

So, we then, also, in the data I’m going to show you, we used the six-month results where we measured the impact on -- measured the success of those who signed up for the account, those who don’t, as well as a control group.

We also, very importantly, will go back after 12 months. And from a science perspective, from a social science perspective, the 12-month results really are the much better results to think about. Why? Because the six-month results, there’s incentives to cheat. There was no surprise factor because it was pre-announced, they knew we were coming in six months to do this, whereas the 12-month results, there was no money on the line at all. Now, we're just seeing whether this continued to succeed in getting people to stop smoking.

Here's a little picture of the urine test. Well, not the urine part, just the test part.

**(Laughter.)**

MR. KARLAN: There was an alternative treatment that we gave people. It was -- by many, you know, it's hard to say what the leading alternative would be, but we wanted to do another treatment that would have really high take-up rates. And there's also a policy that we see implemented in many places. So, in Canada, it is public -- it's law that you have to have these nasty photos on the outside of your cigarettes, you know, the package as you buy it.

And so, we gave out these cue cards that had -- these are the pleasant photos, by the way. The other ones were really much uglier. And they were basically intended to be a cue card that people put in their wallet or in their house somewhere that basically reminds them of the potential negative consequences of smoking. This is meant to mimic the closest we -- you know, kind of a popular public policy.

So, here's the project flow. We start off in the project with a baseline survey. We basically have these bank officers literally just walking through the streets is the exact process. They would walk through the streets to markets, knock on doors of business, and go up to people who were smoking, or even if they were not and just ask them if they were smoking. So, their basic filter was, do you smoke every day? If yes, then they went and filled out a little five to ten-minute survey.

And then in the first two phases of the study, we -- each survey form on the back of it had a sticker assigning people to one of different -- either the CARES treatment group or the CUES cards or control in which nothing more was done.

In the third phase where most of the data comes from, it was randomized through a -- not exactly technically random, but effectively random process by calculating the residual of the day, month, year of their birthdate and dividing by three, and using that to assign. The reason for the change is because we were getting afraid that there was cheating going on in the first method.

So, then they're offered the product. If you'll note in the phase one and two, that we actually had two different CARES products that we were offering, one without and one with deposit collection. Only seven people took up the product without the deposit collection. It was clear that we were not going to have statistical power to separately test out the importance of the deposit collection, so we got rid of it in the full scale-up and only did it with the deposit collection.

So, this now remains kind of a key question for us. It's one thing to have low take-up; it's another thing to have low effectiveness. It would still be very interesting to note what the effectiveness is because you can imagine with a different technology, for instance, cell phone banking, that you might not need the deposit collector, but that's something that remains for us to have to test in a future wave. Then we do the follow-up visit six months or 12 months into it.

AUDIENCE MEMBER: (Off microphone) If they do have deposit collection, they had to go to the bank or --

MR. KARLAN: Exactly, right. So, that's presumably why it was not -- not the most popular option. It wasn't like the bank was that far, but far enough that just too few took up.

So, the probabilities for assignment to treatment did differ across the three ways. In all the statistical analyses we'll show you, we will be controlling for the phase to take this into account. In the first phase, we really were just testing the product concept and the procedures for doing the randomization. So, it was 45, 45 and then just a little token in the CUES and control just to test out the actual procedures.

And then phase two, like I said, we had both deposit collection and not and then we got rid of that for phase three, where we just split things a third, a third, a third.

So, a lot of people don't come back for their urine test, a good number. Enough so that we clearly had to deal with the attrition issue in some way. So, the first thing we do throughout this on the six-month results is we will assume that if you opened up this account and failed to take a urine test, that's because you're still smoking. So, under no circumstances did we assume that the CARES clients who we failed to find or who refused to take the test have stopped. So, they're just assumed to fail.

We want to stack the deck against us, basically, and make sure that our results are not -- you know, not kind of, you know, a result of bad assumptions on attrition.

But then it's not so clear, at least in our minds, what to assume about the people who have no incentive to run. So, we do it a few different ways. In one setting, we'll simply assume that anybody who does not take the follow-up test is still smoking. Another way of dealing with it is perhaps the simplest, which is to assume that they are not different than the ones who are found; hence, we basically just dropped them as observations, which implicitly assumes they have the same quit rate as those who we do find.

So, the specifications -- I'm not going to bother with both. But we do the analysis both with OLS and Probit and find no differences.

There are interesting points to be made about the IV versus the OLS and how to interpret the treatment in this case. So, I will close with a couple comments on that.

So, some basic usage staff to help get a sense on what we have here. So, we have 85 accounts. The number of deposits made into the account, the minimum was one. There was people who opened and never did anything more. The average was 12. The balance at six months, the exchange rate here is 50 to one. So, the balance at six months is, on average, $11, which is pretty significant, actually, I think for this population. The balance at six months of those who returned on average was even higher, above $20, whereas the average for those who forfeited was much lower, thank God.

And then the proportion of clients who missed three deposits -- if you missed three deposits, then the deposit collector stopped coming. It doesn't make any sense any more. And that was only 14 percent. So, for the most part, they -- just tracking usage, it was clear that it was being used for the most part.

Baseline measures. So, this is going to show you a little bit about why we got nervous about the randomization routine that we were using or lack thereof in the phase one and two. And oddly enough, it's the CUES treatment that seems to be mostly off-balance, not the CARES. So, as you can see on the bottom row, for instance, wanting to stop sometime in your life, the CUES treatment is significantly lower than the control, 69 percent versus 75 percent.

So, more interestingly, we do find on who's taking up kind of what you would like to see in the simple correlation results here. So, you know, those who take up are more likely to want to stop smoking than those who do not take up. That's good. That means they're understanding.

Flip to the next page, want to stop smoking now, 29 percent in the CARES group want to stop now whereas only 16 percent of those who do not take up want to stop now. Again, if you look in column six, you'll start seeing where there's some imbalances in the CUES.

So, one of the things we're going to do in all of the analyses I'm going to show you is put in controls for all the covariates. And that's just -- it doesn't actually change any of the results, so we leave that in as our primary specification just out of fear, out of why we may have been finding this imbalance with the CUES in the assignment to treatment.

So, the most interesting result here is on the sophisticated versus naive question, the question that Stephan alluded to at the end of his talk. Who is it that is sophisticated enough to know I am time inconsistent, I need a commitment contract in order to achieve my goal.

So, the closest that we can think of in this setting is asking -- so, it's similar -- I always go back to food analogies for me. For those of you who know me, you know I've been doing contracts like this for about nine years on food. So, I love Indian food. I very rarely go to Indian restaurants. Why? Because I know that I'm going to have a hard time walking in and just ordering the fish tikka and naan, having the naan and all the other things. I would love an Indian restaurant that offered me fish tikka at its normal price and just took the rest of the menu and jacked up the prices really high. I would still love the fish tikka and I'd still enjoy myself and I wouldn't be tempted.

Similarly, in smoking, same basic idea. Are there people who -- the point is, now, I don't even go to Indian restaurants. So, what's the analogy on smoking? The analogy on smoking is if you're in a place that has a -- you know, friends that smoke, if you go to bars where smoking is done, if you want to stop yourself from smoking, then you maybe don't even put yourself in those situations in the first place because you know you'll be tempted. So, you don't go to the bar, you don't hang out with those friends, you don't go to the party where you're going to be exposed to the smoking.

That's exactly one of the things we find that predicts take-up. In the second row there, tries to avoid areas that make him or her want to smoke. Sixty-five percent of those who take up report doing that, whereas only 50 percent of those who don't report that. So, this is the same results just using multi-variate. Let me skip that since I'm running low on time.

Here are the key results. So, these are the intend to treat results using OLS. I'll show you both the six-month and the 12-month. And then, basically, in the different columns what we have are the different assumptions about how we deal with the drop-out, the people who we fail to test. And the results, as you can see, are fairly robust to whatever assumption we put in.

So, you know, what we find is a 3.3 percentage point impact on the -- keep in mind, this is on the intend to treat analysis. So, we have one out of nine taking up. Eleven percent of the people who were offered took up. So, what we notice here is that the CARES treatment is actually not statistically better than the CUES on the intent to treat analysis. Why is this? It’s, hopefully, fairly straightforward. We have 100 percent take-up in the CUES treatment, but only 11 percent in the CARES. So, the intend to treat is diluting the effectiveness of the CARES radically in that sense.

When we look at the treatment on the treated, as you can see, now the CARES treatment blows up to nine X roughly and we have a 30 percentage point improvement in the likelihood that someone stops smoking relative to the control group, whereas the CUES card stays where it was because there was perfect take-up. So, it’s not different. And, now, we have a statistically significant difference between the CARES and the CUES, which is kind of the best analysis to do on an apples to apples basis.

Now, like I mentioned earlier, this is really important to note from a public policy perspective what this is telling us. It’s telling us that if the goal from a public policy perspective is to do the best you can for those who really want to stop, then CARES is clearly better. If your goal is just to treat everyone and figure out what’s the best point estimate for this society as a whole, now we can no longer make an argument based on these results and this sample size that CARES actually did better for everyone as a whole.

Now, you know, clearly, the optimal strategy might be some sort of dual approach. You want to offer things of this nature to those who are sophisticated enough to be aware of this need and to sign up. That doesn’t mean you should ignore everybody else and something has to be done on the others as well.

I’ll just go to this. I’m going to go to this one. Some of the key questions here to ask about, you know, kind of where to go from this are, basically, on how this product design actually works and what are the features of it that are kind of necessary to make this work, how important was the savings box, for instance, and the deposit collection process. And then, also, from a theoretical perspective, better -- you know, we need a deeper understanding of who is it that takes up this type of product, this question of sophisticated versus naive, I think, is very, very important.

The one thing that’s important to note is that the deposit collection and the process of having a bank officer come, in itself, is an element of a commitment contract. So, if this works strictly because of the deposit collection, it’s not a criticism of this being a valid test of time inconsistent preferences. What it means is there’s two things that happened in terms of raising the price of smoking. There’s financial and there’s social shame. And the deposit collector could be working because of the social shame of a bank officer coming to your door and saying, have you stopped this week, are you still smoking, give me your money.

And even if it’s not about the money, there’s a social shame factor and simply opting into a system, into a process that you’re going to have someone come to your door and shame you, that’s, again, still evidence of time inconsistency. It’s evidence that you don’t need the financial contract to do it, you just need social pressure. But that there’s still -- you’re still doing this same exact behavior, still raising the future price of cigarettes. You’re doing it either through social shame or through finances. We can’t separate out those two stories in the study that we did.

That's all. Thank you.

**(Applause.)**

MR. LIST: Thanks, Dean, that was very interesting. Now, anything pressing?

**(No response.)**

MR. LIST: Great. We're going to jump to a quick discussion from Robert Letzler who is a recent grad doing behavioral economics from Berkeley and now he works here at FTC. And I want to publicly thank Robert because he screened all of these papers with me and he was part of the process as well. So, thanks a lot, Robert, for your help.

MR. LETZLER: Thank you. So, the first thing I should observe is I was supposed to have a tag team partner for this discussion, but he couldn't make it. So, I may wrestle with these ideas all by myself, which leads to the next disclaimer that the opinions expressed here are just my own.

So, I think actually our speakers have gone through my first two slides pretty well. So, I will blast through them quickly. But both the commitment device papers and the credit paper contribute to a present-biased literature. My favorite example of present-biased is today I'm thinking about what I should do tomorrow and I say, I really should eat properly and I should go to the gym.  Tomorrow, I eat potato chips and I veg out on the couch and the day after tomorrow I said, really, I should have eaten the broccoli and I should have gone to the gym.

So, as Dean Karlan has pointed out, if you know that you have this problem, you're going to seek commitment devices.  So, these two papers tested two of the plethora of testable implications of this model.  So, Karlan's paper has, I think, very interesting good evidence that some people are sophisticated, they know they have self-control problems, and so, they demand commitment devices and successfully use them.

The next thing I would like to point out is for neoclassical curmudgeons, and I can go into that mode once in a while.  They can kind of explain away some of the other famous commitment device papers.  There is -- famously, Benartzi and Thaler have a save more tomorrow paper.  People commit to saving more money.  This, for a lot of people, is actually committing to a good thing. Then, again, neoclassical out, I can quit any time.  So, it's not really a -- it's a fairly weak commitment.

Ariely and Wertenbroch have something stronger. They say, students in my class, do you want to opt into deadlines.  If you do, if you sign up for a deadline, we professors will enforce them and grade you down if you miss them. But, again, if I have friends who want me to do stuff and give me these social obligations, I need to beg off, it may be very handy to have these deadlines around. I can't do it this weekend, I've got to study.

On the other hand, CARES participants are volunteering to be fined. My sense is if I were teaching undergraduates, I would have an exam question, please explain to me why our theories say you would never do that. If we have actors with no self-control problems, participation may be a dominated strategy if either -- if you can imagine any scenario in which it would be rational to delay quitting or if there's any scenario in which the test could err.

Maybe a neoclassical person could say, I anticipate some future guilt if I forget to pay -- I'm sorry, if I forget to quit smoking, then second-hand -- people will be breathing my second-hand smoke. If I am giving this contribution to the orphanage, I'll feel better. But it's not that satisfying.

On Stephan's paper, as he says, there's lot of evidence as present-biased in the lab. There's plenty of stuff out there that looks and feels like present-biased. Finally, we have this and a few other papers that really have strong measures of lab tests for cash and actual field behavior linked to each other, which is a great thing.

One sort of warning to future authors is we, basically, are identifying everything off of two sets of questions that are measuring time inconsistency. For anyone doing this again, please nail that down better.

James Hilger’s paper, the design and the analysis is convincing and very thorough. My personal bent is I’d like to see this better tied to a big intellectual project. So, what do we learn about how consumers make decisions?

One of the stories, I think James actually didn’t get to talk about was, for high-priced, high-quality wines, there’s no impact of labeling. We’re only seeing it on the low-priced wines. So, are the people who are buying those high-priced wines not responding? Is that because they invested in complementary expertise or is it because they’re overconfident about their wine knowledge.

And, lastly, on Cary Deck’s paper, this is a really fascinating question. Suddenly, we’re moving into a world where mom-and-pop Internet retailers have off-the-shelf access to very sophisticated strategies. In fact, I think most of the time when I’m shopping online, I am buying a product B, which is shipping and handling. And when I go to Amazon, which I presume to be fairly sophisticated, the kind of firm that could hire Peter Crampton or Hal Varian or someone like that to figure out how they should price, they seem to be committing up-front buy enough and we'll give you free shipping and handling. Sign up to be a loyal customer, and we'll give you free shipping and handling.

It also is a pressure -- if I have this fairly fixed sub-additive price of shipping and handling, to either buy nothing or to buy a lot.

In practice, on these Web sites, sort of contrary to the model's assumption, the commitment to buying A when I put it into the cart is, in fact, very weak. I can look at the whole bundle and the shipping and handling and decide to bail out or not.

And the last thing, again, not so -- emphasized in the paper, but not here is it would be really useful to have sort of a numerical solution for the Nash equilibrium which we could use to benchmark what the subjects are doing. Is there time for a few questions?

I think Chris is coming with the cane to get the conference back on schedule.

MR. LIST: Okay, so thanks a lot, Robert, and thanks to all the presenters. I think you guys all did a great job.

**(Paper Session Two concluded.)**

**PAPER SESSION THREE: ECONOMICS OF ANTITRUST**

MR. SHAPIRO: If you could please take your seats, we'd like to start the next session now, please.

Good afternoon. Hi, it's me again. I'm just moderating or chairing this session. The title is economics of antitrust, which means we couldn't quite see how the papers fit together very well, but they're all really good. So, we have, again, three programs. You'll see on your program.

Our first speaker, Michael Waldman, title of paper, Why Tie a Product Consumers Dot No Use? Explanations-Efficiency, Price Discrimination and Exclusion. Michael? Pat DeGraba is going to be the discussant for this.

MR. WALDMAN: Thanks, Carl. So, this paper is co-authored with Dennis Carlton and Joshua Gans. So, this paper is kind of the last paper that Dennis and I, and this one with Joshua, have written as a kind of series of papers on time behavior. Most of the previous analyses -- what we're doing in this paper is trying to put forth a new explanation for why firms might tie. And previous explanations or previous models have typically focused on one of three arguments, either efficiency, price discrimination or exclusionary rationales.

And our argument is a little bit different, which is we have a profit-shifting rationale. The basic idea is suppose you have a monopolistic that ties and the tie itself has some efficiency property to it. What that means is if the consumer just buys the monopolist's primary and complementary goods, the consumer gets a higher gross benefit from consuming the tied good than from consuming individual goods.

And our basic argument is, now add a potential alternative producer of the complementary good who has a superior alternative complementary good by tying, even if the primary -- even if the monopolist complementary good is not going to be used, it serves as an option for the consumer and it reduces the consumer's willingness to pay for the alternative producer's complementary good, and in that way, can shift profits from the alternative producer to the monopolist. That's our basic story. I guess I sort of already went through that slide.

So, let me go through an example just so you can see it in a little bit more of a concrete fashion. Suppose you have Microsoft, which is a monopolist of Windows and a marginal cost, just to keep it nice and simple of zero for Windows, and there's a complementary good and the complementary good is Media Player. So, Microsoft can produce the Media Player and there's a rival that produces QuickTime and the marginal cost of the two is -- the marginal cost for each firm is two for the complementary good.

An assumption which we make which is a little different than most of the literature on the topic is we assume ties are reversible. At least for the example, we're going to have the rival capture all the surplus associated with that superior complementary good.

So, here are the gross benefits. The consumer gets 15 if the consumer consumes Windows and Media Player purchased separately. There's some added functionality if the two goods are tied together, and so, if the consumer consumes Windows and Media Player purchased as a tie, the consumer gets 20. If the consumer -- but QuickTime is a lot better than Media Player. So, if a consumer consumes Windows and QuickTime than the consumer gets 25.

So, if the monopolist sells individual products, if you work through the simple arithmetic of the numbers we just gave, what you get, you get that Microsoft earns 13 per consumer. Why is that? Well, you have this extra ten -- 25 minus 15 that the consumer is willing to pay. Microsoft can get 15 for its product minus the two marginal costs. That gives Microsoft 13 profit per consumer. If Microsoft ties, now it gets 20 per consumer. The consumer's still buying the alternative producer's product. So, now Microsoft's getting 18 per consumer.

So, you're getting this tying, which is the tied product is not actually being used by the consumers, but it's just a shifting of profits. If it costs Microsoft something to produce or put the -- its complementary good onto Windows, then that's socially un-optimal. Or, as we'll talk about later in the presentation, if it costs Microsoft something to produce the alternative good in the first place, then the R&D expenditure is wasted.

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MR. WALDMAN: Excuse me? So, the question is, in this example -- and I'll talk about it a little bit more in the model -- in this example, it's identical consumers. So, there are multiple equilibria and the prices are being chosen simultaneously and the question is, which equilibrium is being chosen? Our assumption here that the rival captures all the surplus is a particular way of picking a particular equilibrium from this multiple equilibria. As I go along, I'll talk about how that comes into play in terms of the general analysis.

In the example, we're making a specific assumption. In the general analysis, we're not going to have that assumption all the time. We'll allow this surplus sharing to be kind of anywhere between zero and one.

AUDIENCE MEMBER: (Off microphone) So, I guess the thing I really don't understand is if they're paying 20 and buying Windows and Media Player and they're buying QuickTime as well --

MR. WALDMAN: That's right.

AUDIENCE MEMBER: (Off microphone) -- why are they (inaudible) -- do they get the extra five surplus (inaudible) Media Player for not using it?

MR. WALDMAN: They're not getting any -- the consumer's not getting any additional surplus for Media Player. It's a question of how much is the consumer willing to pay for QuickTime.

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MR. WALDMAN: So, it's only willing -- so, if I'm owning a good, which if I'm a consumer, I get $20 worth of gross benefit, and if I buy this other good, I get $25 worth of gross benefit, I'm only willing to pay $5 for the good. If I'm owning a good -- if I'm owning a set of goods which -- or potentially owning a set of goods which only give me 15, now I'm willing to pay $10 for the good. So, there's this extra functionality which actually winds up not being used in equilibrium, but it reduces the consumer's willingness to pay for the alternative producer's product and that's going to help the monopolist at the end of the day by shifting profits from the alternative producer to the monopolist.

AUDIENCE MEMBER: (Off microphone) Why doesn't -- just (inaudible) it will help (inaudible) follow up. Why doesn't Microsoft just charge $20 for Windows without tying it (inaudible) consumer doesn't buy it (inaudible).

MR. WALDMAN: Well, again --

AUDIENCE MEMBER: (Off microphone) Does that make it -- and QuickTime -- I can see an equilibrium when Microsoft charges $20 for Windows; Apple or whoever it is charges $5 for QuickTime. That's equilibrium.

MR. WALDMAN: Right. So, the answer is if we assumed sequential price-setting, then the result goes away because you could do exactly what Carl just suggested. What we assume is that the prices are set simultaneously, in which case, there's a multiple equilibria problem and then the question is, how do you resolve that multiple equilibria problem? Depending on how it's resolved, you can get a return to time, which is basically the whole argument.

Does that help you out?

AUDIENCE MEMBER: Yes.

MR. WALDMAN: Okay. Now you know the whole paper, but anyway, I'll go through the rest of the presentation anyway.

**(Laughter.)**

MR. WALDMAN: So, this paper captures and extends the logic of that example. We have a monopolist of a primary product, a complementary good that can be produced by the monopolist and a rival. Consumers only have valuations for systems, just like in the example we just did, and ties are reversible.

I'll go through or the paper goes through three different analyses. It goes through an identical consumer analysis along the lines of the example we just did. That's a simple case of heterogeneous consumers and also has our endogenous R&D choice by the monopolist. In each case, what we show is that you can get tying when, at the end of the day, the tied product, the complementary good that's being added to the primary product, is not used in equilibrium. So, that's a social welfare loss if there's some cost of producing this complementary good. And we also have a second social welfare loss which we talk about in the third model, which is this R&D cost of producing this complementary good in the first place.

So, the talk is going to be, as much as I can get through, relationship to the previous literature, model analysis, R&D distortions, just a brief discussion on antitrust, since this is labeled antitrust, and a conclusion.

So, if you look at most of the previous models, and I've listed a few of them, Whinston, Choi and Stefanadis, an earlier paper I did with Dennis, and Nalebuff where the tying is used to disadvantage rivals, mostly the tying is used either to cause exit or block entry.

And what we're doing is quite different than that. So, our tying is used to disadvantage rival, reduce the profit of a rival, but the goal is not to stop the rival from entering the market or cause the rival to exit the market, rather the goal is to take some of the rival's profits and shift it over to the primary good producer.

Now, a famous result or an important result in the tying literature is a result that goes back to Whinston which shows if the monopolist primary good is essential, then there's no return to tying. What do I mean by essential? Essential means that the monopolist -- that the monopolist's primary product is used -- if you want to use the complementary good, you have to get the monopolist primary good. And in Whinston's models, if that was the case, then there was never a return to tying.

What we show, in this alternative model, is that you can get a return to tying even if you have this essential element. Our model's a little bit different. A, we assume that ties are irreversible -- I'm sorry, Whinston assumed ties are irreversible and no efficiency associated with the tie. What we're showing is that if you remove those two assumptions, irreversible means if I tie the product, I can't add someone else's product to the tied system. And, so, especially if I'm talking about Microsoft, that's not that realistic. So, we more realistically assume reversibility and we also assume this possibility for efficiency associated with the tie and when you allow for those two things, then Whinston's essential result winds up going away.

There's also a couple of papers that look at independent products where the role of the tie is to reduce competition in one of the markets, Carbajo, de Meza and Seidman. Our paper is a little different than that because we're allowing -- we have complementary goods and the role of the tie is quite different than what's going on in those two previous papers.

Actually, our paper is closest to a paper by Farrell and Katz, which was in the Journal of Industrial Economics in 2000. What they basically show is that there's various behaviors that a firm might take on to create a price squeeze for the completely good producers, integration, R&D, exclusionary deals. What we're showing is that basically their argument also applies to tying if you allow for reversible ties and you allow for these efficiencies associated with the tie.

So, here's the model. I'll go through it very quickly. We have a monopolist and a single-alternative producer in a one-period setting. There's a constant marginal cost for the monopolist for the primary good of CP. The complementary good, both for the monopolist and the alternative producer, there's a constant marginal cost of CC and the alternative producer's complementary good is superior.

Goods are only consumed in systems. Ties are reversible and tying is weakly efficient. So, basically, if the consumers consume the alternative -- the monopolist tied primary and complementary goods, as opposed to the alternative producer's complementary good, there's at least -- it's weakly efficient to have them tied. And in this first model, it's identical consumers.

Here are the gross benefits, VM if an individual purchases the monopolist's products and consumes the monopolist's products purchased separately. The consumer gets an extra delta if the product is tied.

The consumer gets VA if the consumer has the monopolist's primary good and the alternative producer's complementary good, and it's then the max of VM plus delta and VA if the consumer buys the monopolist tied product and then adds the alternative producer's complementary good.

The timing of the game is the monopolist decides whether or not to tie. Oh, just one quick aside, the current version of the paper is a little bit of a mistake. We claim that it generalizes the mixed bundling. That's not actually true and we're in the process of rewriting the paper to fix that up.

Firms choose prices and then consumers make their purchase decisions. We look at a sub-gain perfect Nash equilibria, and as I was saying earlier, this is a -- it's well-known or at least people who work in this area know that this is a gain where there's frequently multiple equilibrium and that's true of this gain, and we're going to resolve that multiple equilibria problem by assuming that -- we're going to assume that the alternative producer's product is superior, the monopolist gets -- I'm sorry, the alternative producer gets lambda of its superior product and the monopolist gets one minus lambda.

If I have time, I'll actually talk a little bit about what happens when you move away from that strong assumption and allow lambda to vary whether or not the monopolist ties or not.

So, parameter restrictions, we assume that it’s efficient for the monopolist to release its products, which means VM greater than CP plus CC, delta is greater than zero, greater or equal to zero.  That just means there’s an efficiency -- weakly efficiency associated with tying.  VA greater than VM means the alternative producer’s product is superior.

So, the first result is to basically generalize Whinston’s result, which says, if there’s no efficiency associated with the tie, if delta equals zero, then there’s no reason to tie in this world.  So, Whinston’s result goes generalize, even with reversible ties in our model, as long as there’s no efficiency extra functionality associated with the tie.

But if we allow this extra functionality, delta greater than zero, then you get the following set of parameter values translating to different types of behavior, and you can take those five regions of the parameter space and translate them into efficient and inefficient behavior.

So, parameter condition one translates into efficient tying; three and five on the previous slide translates into efficient sales of individual products by the monopolist; and then two and four translates to inefficient tying. So, four, in particular, is basically just the generalization of the example that we did in the -- that I did in the introduction which is, the monopolist ties, the consumers go ahead and purchase the alternative producer’s product, don’t use the alternative producer’s product, but, in fact -- and, so, there’s a social welfare loss there. But that’s a way for the monopolist to increase its profits.

There’s also a part of the parameter space captured and, two, where the monopolist ties inefficiently, but actually the consumers go ahead and use the monopolist product.

We can change the sharing rule assumption. If we had lambda equal to zero, you’re still going to get similar results where you’re going to get inefficient tying. You could also have lambda be greater than zero when the monopolist sells individual products and lambda equals zero when the monopolist ties and -- or in other words, allow for different values for lambda with tying and without and you’re still going to get different values for the parameter ranges, but you’re still going to get two parameter ranges with inefficient tying just like I had on the previous slide.

We also have a short analysis considering heterogeneous consumers just showing if you introduce a small number of -- a second group who have different preference characteristics, you can still get similar results.

Then what we do -- in terms of thinking about this as Microsoft, you might say, well, gee, I'm not sure that the Microsoft's marginal cost for having a -- putting its goods onto Windows is really very high, and so, that might be a social welfare loss that I really don't care about it. It might be second order.

So, we spend a few pages talking about, well, suppose we add R&D decisions into the paper, and so, in particular, we allow this delta, the extra functionality to be either small or large where it's a function of the amount of investment that Microsoft makes into the R&D process for producing this complimentary good. And then we stick with the same parameter range from the previous analysis where you got this inefficient time where the good wasn't actually used but still purchased.

And to make a long story short, you can go through that analysis and what you find is you get a second social welfare distortion. What's the second social welfare distortion, the second social welfare distortion is that Microsoft winds up investing in this R&D even though, in this part of the parameter space, the consumers never actually use Microsoft’s products. So, Microsoft frequently ties and the consumers aren’t using the product and, yet, it’s -- the reason it’s tying is because it’s a way of causing, an expected value sense, more of the profits to be shifted from the rival to the monopolist.

And then we also do -- this is not in the paper. We also do a second analysis where the alternative producer has an R&D decision and what you find there is that this type of behavior can cause a distortion by causing the alternative producer to invest less in R&D.

Just very quickly on antitrust, Dennis and I have a couple of papers, one coming out in Antitrust Bulletin with Patrick Greenlee, talking about antitrust -- our views of antitrust here. So, our general views for antitrust are in those other papers. In terms of the second paper, we talk about safe harbors and that they should be based on main theories concerning the harmful effects of time.

In terms of the current paper, what we would say in terms of antitrust is one should be very careful about using this to sort of stop time because you only see this effect when there’s an efficiency associated with the tie in the absence of the alternative producer.

So, it would be very hard for the Courts to really figure out that this is really what's going on as opposed to it's just a standard efficiency argument where, in fact, Microsoft just -- or the monopolist just wound up not producing a good enough product.

Conclusion. So, what we've done in this paper is provided a new explanation for tying, which is basically a profit or a rent-shifting explanation. And in terms of kind of focusing what -- more specifically what have we done, we realistically allow time to be reversible which is, I think, an advancement over the previous literature. We show why a firm might tie even if the consumer's product is not used in equilibrium and that seems consistent or at least roughly consistent with some of the things Microsoft tends to do.

We show that Whinston's result concerning the essential nature of the product be important as to whether or not you see tying is not robust to this. And, finally, from an antitrust perspective, what we would say is we think that in this particular theory for harmful tying, one should be very careful in terms of using this as a basis for antitrust intervention just because it only works when there's this efficiency. So, it would be very hard for the courts to pull out that this is exactly what's going on as opposed to that Microsoft or the firm was trying to do something efficient and that wasn't quite as successful in terms of the quality of the product that was produced.

MR. DeGRABA: So, it's an idea that I like, but it's a model that I think could be streamlined and could have avoided most of the conversation here about multiple equilibria. In fact, this paper -- 90 percent of this paper, actually, isn't about tying at all. It's more about, sort of, the Cournot Complements problem and the idea is simply that if one firm has a primary good and a competitor is offering a complementary good, a really high price for the complementary good means the monopolist can't sell the primary good at a very high price or sell very many of them.

Anything you can do to get the complementary good provider to lower the price means he can either sell more units or raise the price or typically both of the primary good. So, in this paper, if the primary good supplier also offers another complementary good, even if it's inferior, it provides some competition for the good complementary good. It lowers the price of the good complementary good and increases the price that can be charged for the primary good.

In this particular paper, one of the ways to make the primary good -- to lower the price of the complimentary good is to simply tie a product for which consumers have some utility for the tying. That makes the bundle sort of more beneficial, it lowers the price that the superior complimentary good can charge and it raises the profits of the firm.

What I want to talk about probably will bore most people, but probably not Mike. And I want to argue or I want to at least suggest, I spent a week thinking about the paper, and I want to simply say that if you -- instead of having a model where all the consumers value the primary bundle the same, if you have some price sensitivity to that, you can get rid of the multiple equilibria problem. You can get a unique equilibrium and a unique price and you can actually get more results out of that model than you can out of what's in your paper.

So, in the -- I have the new model here, the model I'm going to propose. Instead of everybody valuing the complementary bundle at some constant amount, we're just going to have a linear demand curve for the constant -- for the monopolist bundle. Everybody values the alternate good. That's the same extra amount, E. So, that part is still the same as in the model. And tying still generates some benefit delta. All the marginal costs here are zero.

I'm going to argue that in this paper the dominant strategy for the monopolist -- and, so, the model is then everybody sets prices and customers buy what's best for them. I'm going to argue that the dominant strategy for the monopolist is to set his complementary good at zero and that the alternative complementary good provider would never set the price of his good more than E above the price of the monopolist complementary good if there's no tying and no more than E minus delta if there is tying.

The graph of the -- the incentives to tie are easily shown in this graph. If there's no tying, alternate -- the complementary good producer simply charges E, the extra value of his good. The monopolist sets the price of zero for his complementary good, and then he's basically faced with a residual demand curve, which is just a demand for the monopoly good.

So, the three curves up here, the green one, which is the lowest one, is just a demand for the monopolist bundle. The red one on top is the demand for the bundle with the alternate -- with the good -- extra good in it.

In the no-tying equilibrium, the monopolist sets the monopoly price for his own bundle and the alternate firm captures E for all the units that he sells. When the tying occurs, what happens is that the value that the alternate producer can charge is now his extra value minus the new delta that got introduced. So, he has to lower the price of his good by delta. That shifts the demand curve for the monopoly good up to W plus delta, which is the purple line. That new higher demand curve allows the monopolist to raise his price. So, in equilibrium you have -- the tying prices are E minus delta and W plus delta over two.

Why is that interesting? Well, in this particular case, if you notice the price of the tying good went down by more than the price of the primary good went up. So, this says that if you have this efficiency going on, that tying will actually lower the price to consumers of the overall bundle -- that’s a result that’s not in the paper.

There’s a second equilibrium that I’m not going to go through, but if E is big enough, then it turns out that the price of the ultimate complementary good isn’t really constrained by the monopolist complementary good, that he’d actually rather charge a price lower than E in equilibrium. And in that case, you’ll find that tying won’t do anything at all.

And there’s one other result and then I’ll finish, which is -- which I’m only sort of convinced of because I haven’t actually done the math, which is that in equilibrium, when there’s no tying and there’s positive marginal cost, the monopolist still sets his complementary good price at zero. So, if you were to do something like tell Microsoft not to bundle, say, Windows and some media player, for instance -- and by the way, I’m invoking the disclaimer for FTC employees at this point -- that this model predicts that in equilibrium, Microsoft will give away Windows Media Player and not charge a positive price for it.

The final thing I want to say is that this model had sort of valuations differing for the monopolist good and everybody viewed the alternate good with the same valuation. If you reverse those assumptions, the model becomes much more complicated, but also, I think, much more interesting. It’s something that ought to be worked on. Thanks.

MR. SHAPIRO: All right. Let’s move right along. So, as you can see, our next speaker is Minjae Song, Sleeping with the Enemy: Inter-firm Product Combinations.

MR. SONG: This paper is with Claudio Lucarelli and Sean Nicholson. Both of them are at Cornell. The paper is still preliminary.

So, this paper is about the product combination, but inter-firm product combination. So, it’s not the combination of the products produced by the same firm, but by the different firms. And this practice is kind of prevalent in the pharmaceutical markets. For example, in the AIDS treatment or the chemotherapy treatment for breast cancer, lung cancer or colon cancer. And in this paper, we’ll look at the colon cancer market.

So, the prescribed, we call it, regimens consist of two or three drugs produced by the competing company or different companies in the market. And these regimens are not bundled goods, because physicians buy individual drugs in the market and then they make cocktails and they prescribe regimens, following sort of recipes provided by National Comprehensive Cancer Network.

So, let me give you a few examples of these cocktails in the market. The first one, there’s a stand-alone regimen produced by Pfizer and the second one is the stand-alone regimen by Roche. But the third one is a cocktail. It consists of a drug by Roche and Pfizer. You can see that the dosage in the third one is slightly different from the first two.

So, in the first one, you have 125-milligram per week for four weeks to every six weeks. But in the third one, it’s the 70-milligram every six weeks and also the dosage for the Roche -- the drug also goes down from

2,500 to 2,000. And the first example is the stand-alone -- the regimen by Sanofi. This one is another cocktail made of the drug by Sanofi and then Roche.

So, we treat each -- the recipe as a product. And we have the market share data on each regimen in the market.

So, the question we have here is whether these cocktails increase the firm's profits, and if it does, then how much it would increase the profit. And whether the cocktails make drug prices higher or lower compared to the market without a cocktail.

We are also interested in consumer welfare. So, we ask if consumers are better off or worse off with this cocktail. So, if you take out this cocktail, then -- I mean, if you provide this cocktail, there's more variety the wider set of the choice, so a consumer may get better off with a cocktail, but there's also a pricing issue. So, depending on how the price changes, with a cocktail, the consumer welfare may get affected.

So, our approach is kind of an empirical approach since the -- kind of the cocktail structure in the market is quite complicated, we look at data and then we estimate regimen level demand model using the discreet choice model. Then we assume that market is Bertrand Nash equilibrium. So, the firms set drug prices and the price that we observe in the market is in Bertrand Nash equilibrium. So, we can recover the marginal cost from the estimated demand.

Then given the demand estimates and the marginal costs, we did two counterfactual exercises. The first one is we take one cocktail out of the market one at a time and then we recompute the equilibrium price and then see how the profits and consumer welfare change.

The second counterfactual -- in practice, in the market, the firms can only set one price per drug. But in the counterfactual, we allow a firm to set two prices. So, one price for a stand-alone regimen and the other price for the drug used in the cocktail regimens. It’s kind of an interesting exercise because in the AIDS treatment segment, there’s a company called Abbott and they actually have two drugs. They used to have only one drug and that drug was used in a cocktail, combined with its rival company. Then they were about to launch a new drug in the market, and to secure the market share for this new drug, they increased the price of their existing drug by about five times.

So, there was an article in the Wall Street Journal, and I think FTC is kind of looking at that case this year.

So, we kind of tried to mimic that case, although we don’t observe that in this market. So, let’s think about this problem -- is there a question or --

AUDIENCE MEMBER: Yeah. In looking at prices (inaudible) are you looking, however, at (inaudible)?

MR. SONG: No, that’s something that we cannot control. So, we are only looking at price and profits and consumer welfare. So, for example -- so, when we take out -- the regimen out of the market, this attribute issue doesn’t exist there. But when we allow the firms to have two drugs, there’s nothing we can do about those hypothetical attributes. So, we just make an assumption that the attribute doesn’t change, efficacy doesn’t change, but they only allow two set prices. So, we are not looking at the efficacy or the side effect issue.

So, let’s think about this problem in a very, very simple setting. So, there are two pharmaceutical companies in the market and each one has one drug per each. So, you can have a regimen one, regimen two, which are stand-alone regimens. And regimen three is a cocktail regimen, so it consists of the drug one and drug two. And Q11, Q22, Q13, Q23, so dosage of each drug used in each regimen.

So, you can write down very simple profit functions say for firm one and the profit comes from the two regimens. One from the stand-alone, what we call the solo regimen, and the second source of profit is this cocktail regimen. And assuming the Bertrand Nash, you can solve this profit maximization.

The first little condition is slightly different from kind of the standard one that we know because of there's two abstract (inaudible) here because you not only care about the elasticity, the effect of the share from the price change of P1. It's only elasticity and cross-elasticity. But you also care about the effect on share through the price of the third regimen, which is the function of the two prices here.

So, the stand-alone regimen is just how much dosage you use times the price per milligram. But the price for the cocktail regimen is the function of your price and your rival's -- the price. And when you set the price one, you have to look at the effects through the regimen price -- the price of regimen one and the price of regimen three.

So, this is kind of the simplest setting that we can think of. In the paper, we did some numerical simulations, given this very simple duopoly setting. We do the kind of two counterfactual data. I'm going to show in this simplest setting and then we kind of see -- we analyze what happens. It's not really the same -- we don't get the same jumps in the data because of the more complicated -- this cocktail structure. But there's something that we can learn from this simple setting, but I'm not going to present the numerical simulation here.

So, demand estimation is very, very standard, the discreet choice demand model. We start with the loaded model (inaudible) model with our random coefficients. We may try the random coefficient (inaudible) or even the pure characteristics demand model. The reason is so when you write down this (inaudible) model, you know, you have to define, you know, who the consumer is. So, we kind of assume that the physicians are kind of imperfect agents for their patients. So, the product attributes and price part of the alternative function, the physicians care about their part because they want to cure their patient with the most effective drug or the regimen.

But there's also (inaudible) in the logic demand model and we kind of interpret that as some -- the inconsistency between the patient's utility and (inaudible). And one example of that is kind of the rebate that the physicians get from pharmaceutical companies at the end of every year. So, this is kind of the current model that we have for demand. But, for example, if you move to a pure characteristics demand model, there's no idiosyncratic shock. So, the physicians care about the patient's cure. So, that's kind of an alternative demand model that you can think of. This is a very standard model.

So, there are 12 regimens in the market and the six of them are cocktails. So, there are examples of the cocktails. Pfizer has a cocktail with Roche, Pfizer has a cocktail with Genentech, Pfizer has a cocktail with Imclone and also there are other cocktails in the market. I'm not going to show you to go on.

And we have drug price and then regimen market share, and based on this cocktail recipe that we get from NCCN, we can compute the regimen prices using the drug prices and the cocktails with price and market share at the regimen level and we have regimen attributes based on the three clinical trials. And they include efficacy like survival rates, time to progression and two more response rates. We also include the side effects, like vomiting, diarrhea and et cetera.

And then demand estimation is very, very straightforward. It's just the simple logic model, but we have to control the price endogeneity. So, we need the IB. So, we first tried with the standard PRB IB, but that's really a good idea because the characteristics do not change over time. So, a characteristic change only when you have new products or you've just seen product exit. But we don’t observe any major exit here.

So, the first stage statistic of the (inaudible) IB (inaudible) 16. So, the second type of IB we tried is we used the other regimens’ prices in T minus one. So, the assumption is the prices also correlated over time, but demand shock is not. So, this is kind of our identifying assumption.

We’re concerned about the weak instruments, but (inaudible) is over 60 on the first stage. And this is our (inaudible) and IB and (inaudible). What is interesting is even without any instruments, we have negative price coefficients. So, it’s actually physicians who make this decision and they are about price. And the first IB doesn’t really change much of the coefficients. But the second one, it decreases the price coefficient a lot. Some of the signs are not really interesting. So, the time to progression should have a positive coefficient, but we have a negative.

So, the only -- I mean, in all specification, the only coefficient that makes sense is response rate. But what we think is the -- is how to separate the three efficacy from each other. So, when you kind of think of the utility from the drug, you have to think about it as a combination of this efficacy.

And side effects do not really come out in a significant way mainly because physicians also control the side effects by giving like a drug for the diarrhea and a drug for the vomiting. We think that could be the way.

So, let me -- in my remaining four minutes -- three minutes, okay. Let me show you my counterfactual. So, the first counterfactual, we take out one regimen at a time. There are six cocktails in the market. So, each row corresponds to each exercise. So, for example, the current one is the one that we observe in the market and we normalize it at 100. So, everything is compared to what we observe in the market.

The second column is the case where we take Pfizer and Roche out of the market, and price of Pfizer, this is like a single drug, goes up by like 14 percent or 15 percent, but the Roche's, the price goes down about 9 percent. So, if you kind of look through this bold letter face, there's no kind of clear direction of the price changes. Although in the two-firm duopoly, very, very simple numerical setting that I kind of briefly mentioned, what we observe in that numerical exercise is when you take out a cocktail, the price actually goes up, conditional on both firms getting benefit from this cocktail. But this more complicated structure of the cocktail doesn't result in that kind of results.

This is the profit. So, the profit shows that for every case, the firms get worse off without a cocktail, without a cocktail. So, some of the profit changes are very, very significant. For example, in here, Imclone's profit level without a cocktail is about 21 percent of its current -- the profit. The reason is that Imclone's cocktail has larger market share than its stand-alone.

You look unhappy, so I will finish. So, Imclone's -- the cocktail's market share is much higher than its stand-alone cocktail. Another example is Genentech. The profit goes down to the 25 percent level of the current level because Genentech doesn't have any solo regimens, stand-alone regimens. All of Genentech's products are the cocktail regimens. So, they get hurt a lot by the -- without the cocktail.

Consumer welfare. So, consumers care about two things. One is the variety and second is the price. So, when you take out cocktail, the welfare goes down because there's one less product on the market. But in two cases, they're actually better off without a cocktail. Why? Because without those cocktails, the price of other drugs goes down so that they get more benefit from lower price than extra products in the market. So, we have kind of an interesting result.

So, the second counterfactual -- give me one minute and I'll kind of finish. The second counterfactual is with allow a firm to set two prices, two separate prices. So, the first column here is the price for the solo drug or the stand-alone regimen. So, let me look at this. The Roche case, when Roche set two prices, they lower their stand-alone price to the 67 percent, but they increase their drug price for the cocktails about five times.

And in Sanofi's case, they lower their price for the solo regiment about 27 percent, but they increased their drug price for the cocktail about twice, which is kind of consistent with what we observe in the AIDS market.

Interestingly, this practice doesn't necessarily increase the profits of the pharmaceutical companies and we kind of talk about why that could be in the paper. So, since I ran out of time, I'll just finish here.

MR. SHAPIRO: Thank you. Our discussant is Wei Tan.

MR. TAN: So, first of all, I think this is a very good paper. Minjae has done a wonderful job in terms of compounding a very rich data set of pharmaceutical demand. These are also very interesting questions, especially to study the pricing strategies when firms use -- may have inter-firm product combinations.

So, let’s just first give a brief summary about the main finding of the paper. So, this paper tried to study the pricing strategies when firms use inter-product combinations. And the strategy they use is try to estimate the amount of systems at the regimen level. And then they tried to recover the cost parameters from the Nash equilibrium conditions. After that, they will be able to perform counterfactuals to evaluate the impact of product bundling.

So, here are some of the comments. So, first of all, it is very important to get the demand estimates right. Because the starting point of your analysis, you have to back out the marginal cost and the counterfactual analysis.

However, right now, the paper uses (inaudible) logic demand functions by basically transforming the market shares and then it becomes a linear function so you can use the IB approach to estimate the price coefficient. The main problem with IB (inaudible) demand function is that the estimates of the demand elasticity are completely driven by the market share of the regimens and this is a well-recognized problem using the simple IB (inaudible) demand functions. So, you could try to, for example, adding random coefficients and also especially the additional consumer characteristics and also the interaction in terms of consumer characteristics with the (inaudible) observed (inaudible) observed the random coefficients, and that's shown in the literature to be very important to get the demand elasticity right.

Additional robust check, the market structure changes a lot during the sample period. In the beginning of your period, there is one drug that has almost like a 90 percent of market share, it's like a monopoly. And at the end of your sample period, for these particular regimens, the market share is only something like 10 percent and you have 12 regimen combinations. So, I would say that maybe try to use the later periods, probably you don't have enough observations, I'm not sure, to see whether the estimate still holds or whether you have some changes in your results.

Another recommendation is to do reported estimates of the delta coefficient that you get because this is going to be important. Depending on the value of delta so that will affect the mixing strategies. Also, it would be interesting to know what is the estimated delta coefficients and how that relates to your other simulation results.

Additional comments, so the simulation results is still fairly restrictive of function of firms. I'm just quite curious to know whether your simulation results hold for a more general demand function or it's something like a -- the result of treatment by the peculiar logic demand functions.

Another important thing is the law of advertising when you have product bundling. This is important because your simulation counterfactuals show that when firms choose different prices for a solo regimen and bundled regimen, then you have dramatically different prices. So, you could think about it, if this is indeed true, then this actually makes advertising even more important. You could choose a same price for the product, but you can adjust the demand by changing your advertising intensities for a different regimen.

The differing data probably would be difficult to get, but I think one way to do that would be to try to get some sort of like direct-to-consumer advertising data. This type of data, I think, is usually available and you could get, for example, how much they spend on the direct-to-consumer advertising and for different regimens and to see whether that has any impact on demand functions.

The last thing I want to say is about the bundling and entry process. So, it’s very interesting. For this particular market, the bundling decisions are made by the entrant, not by the incumbent firm. So, when -- so, this naturally raises a question that why do entrants choose to bundle with some incumbent firms but not the other incumbent firms. This also has to do with your counterfactual simulations because it doesn’t make sense to remove one of the regimens that is already there. So, one of your findings is that you have some inconsistent findings about prices. But if you look at only the new entrant strategy -- actually, if you look at it, it is consistent with those theoretical models. So, maybe you should look at the new entrant decisions.

Another thing is that when you look at the -- the entrants’ decisions about whether to bundle with a particular incumbent firm, you have to worry about whether they choose not to bundle with a particular incumbent firm because it’s not profitable or whether it’s not medically effective enough. So, this type of information I assume you can gather from the clinical trial data. You can see whether they have tried to combine with some incumbent firm’s product or not.

The last thing I want to talk about is you could have looked at what is the impact of bundling on the incumbents’ pricing strategies, whether you have an entry accommodation or an entry deterrent story that’s going on here.

To summarize, I think this paper is really interesting and I really enjoyed it.

MR. SHAPIRO: Thank you. Again, we could take a couple questions. Yes, would you step to the microphone, please, so you can be heard. Minjae, you might want to come up here to respond.

AUDIENCE MEMBER: So, my question is whether we should consider dosage to be endogenous at the firm level. I would imagine -- this is out of my naive knowledge of pharmaceutical -- is that if you take out a cocktail bundle, the firm could re-choose their dosage so that consumers can bundle themselves in their own kitchen instead of --

MR. SONG: So, in this case, physicians are the ones who make these cocktails. So, in the paper we take the dosage as exogenous and a fixed amount. But in practice, physicians actually try with a different dosage. So, if this regimen doesn’t work on some of their patients, they may increase some of those. So, in practice, actually endogenous, but then we can do another kind of derivative. When the price changes, how does that affect the physician’s choice of the dosage. But they should be within some boundary because they cannot just change it from 1,000 milligrams to 2,000 milligrams. It should be bounded. But, yeah, that can be potentially (inaudible).

MR. SHAPIRO: (Inaudible).

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MR. SONG: We are planning two more things. So, we will get -- not me, but my co-authors will probably get the physician level data because the market share data was computed based on the physician level data. So, if we get physician level data, then we may be able to model the physician’s choice, like, you know, the way they use (inaudible) it’s kind of matched between the drug and the physician and there’s some risk of (inaudible) in this match. Then we can probably come up with the richer model of the demand.

Right now, we have only the regimen level data, so we can still kind of stretch our model to accommodate those features, but -- yeah, uncertainty. Uncertainty, yes. I would kind of think about that, you know, on how to -- how we can do this with data.

And the second thing that we are going to do is we have the clinical trial level data. That’s kind of not really related to your question, but to our discussant’s comment. So, we kind of tried to link the decisions in the clinical trial phase to the -- what’s going on in the product market. But this uncertainty and risk of averseness is very important. Thank you.

MR. SHAPIRO: Well, let's stop there, thank you. Our third paper here is by Christian Rojas, the Role of Information and Monitoring on Collusion.

MR. ROJAS: Okay. So, this paper is experimental in nature, so I was not in the experimental stages. So, I hope I do not disappoint some of you by presenting some experimental results.

I should put the work in some context so that you get a little bit better focus on what I'm going to talk about. The context of the paper is on repeated gains with demand uncertainty. And, so, the information that I'm talking about here is going to be the information related to how much firms know about the demand schedule they're going to face in the future. And monitoring obviously relates to what firms know other firms are doing in the market. So, that's the context of the paper.

So, let me talk a little bit about the motivation. I think this paper, in my view, is very ambitious because the different things I want to tackle. So, I think it's really important to talk about this.

I think there are two stylized facts in IO about the factors that affect collusion. One is monitoring. If you're able to monitor other rivals in the market, then collusion should be easier. Another thing at least that some people agree on is that demand information also matters. If you know what demand schedule you're going to face, then collusion should be easier to achieve. So, I'm going to take a look at whether experimentally this is the case and I think that -- I have the second motivation a theoretical motivation, although in the back of my mind I think this is the main motivation, which is the opposing predictions of two well-known theories about cartel stability.

On the one hand, we all know that the predictions of Green and Porter, who basically say -- this theory is known as a theory that predicts finite price wars that are triggered by low demand for a finite period of time. I think I already said that.

In other words, most people know this theory as in which collusion is more stable during periods of high demand. On the other hand, we have a theory by Rotemberg and Saloner who basically predict that price wars should be observed demand is high and collusion, as a consequence, is more stable during periods of low demand. So, of course, there are different assumptions in the models and these assumptions are going to be informing me of the experimental design. But I think studying these two theories is really important because of the opposing predictions that they have and, also, because as many theories of dynamic gains, we have several equilibria.

And an interesting question, at least for me and for some people is, how plausible are these equilibria or are these predictions by these models, for example, as opposed to other predictions that are also equilibria.

So, the last motivation, as you see, experiments to explore this question, and I think experiments can be a useful tool, especially here because we have multiplicity of equilibria. I want to see how likely some predictions are versus others, and also because collusion is an illegal matter here in the U.S., data's really difficult to get. So, that's my general motivation.

So, let me talk about the general set of assumptions that I have for the two theories. You can obviously cast these two theories slightly differently, but you can cast them in a basic set of assumptions and they're going to differ on their assumption about the nature of demand. So, we're going to assume that both theories have homogeneous products, competition is Cournot, firms are symmetric and they have constant marginal costs. This is always (inaudible).

And the key thing here, the key difference between the two models is that Rotemberg and Saloner, although in their setup, demand is stochastic. They wake up and they know what demand is going to be tomorrow. So, some of the uncertainty is removed and that's why I have in both letters that they have perfect information on T plus one, although they don't know what they're going to face T plus two.

The other thing that they assume is that firms have perfect monitoring available to them. In other words, they wake up with a profit and they also know what other people chose as quantity.

Green and Porter, on the other hand, have a much more uncertain environment in which there's uncertainty about all future and past demand shocks and the additional assumption they have is that there's imperfect monitoring. In other words, you wake up with a profit. You see an imperfect signal and you don't know what other people chose in the market.

So, this is basically the difference between the two models and this is -- these differences are going to allow me to construct an experimental design.

So, very briefly, the reason why the Rotemberg and Saloner prediction comes up is because since you know tomorrow you're going to face a big demand shock, then the incentive to collude, which is the left-hand side of that equation, gets really big compared to what the punishment you would get if you actually deviate.

So, when the demand shock is very low, then you have the second equation where the inequality is reversed and you actually have collusion -- you actually have -- the deviation doesn't pay off. Of course, there are other equilibria. I'm going to compare the results that I get with respect to other equilibria.

Very briefly, again, the tradition of Green and Porter is one in which you have an imperfect signal of what's going on in the market. In other words, you can't see what other people are actually seeing as quantities. So, you wake up with a low profit. This could either be caused by a rival's defection or by low demand. You really don't know.

But the equilibria is what I like to call the mafia-like equilibrium where everyone is suspicious about everyone else and there's this imperfect monitoring device, which is price. And if this price falls below a threshold level than everyone starts a price war, even though no one deviated. And this is the kind of equilibrium that they entertain is one in which no one really deviates from the equilibrium -- from the collusion, but they do start a price war when they see a low price. And the only way they can see a low price is because of low demand shock. So, this is the main thing -- the main prediction.

Now, something really important about Green and Porter is it is known as a theory of finite price wars, where you calibrate the N, that N star that you see there, so that you just offset the incentives to deviate. So, you just make firms indifferent from deviating today with respect to colluding. But this also means that if N star would give you an equilibrium, N star plus one, N star plus two, and N equals to infinity is also an equilibrium of the type entertained by Green and Porter. So, there is a multiplicity of equilibria, although the Green and Porter paper is known for a finite price war. So, this is really important.

So, in my experimental design, I have a very simple setup. The reason why I have this repeated prisoners dilemma game is because I really want to give the two theories the best chance of occurrence. I want to give collusion the best chance of occurrence. So, I have two quantity choices, low and high. Low obviously is the collusive outcome or the collusive choice. And I have three demand states. A very simple, stochastic demand schedule where I have high, medium and low demand with the probabilities that you see there.

This is a repeated game. The way I deal with the infinitely repeated aspect of it is using a procedure that experimentalists have used in the past, which is a random (inaudible) rule, in this case, 20 percent probability would simulate a discount factor of delta equal to .75.

In the three treatments that I have, the first and the third one that you see there are supposed to resemble the assumptions of the two theories. So, remember on the one stream we have Rotemberg and Saloner, perfect demand information and perfect monitoring, and Green and Porter is at the other extreme where you have imperfect monitoring and imperfect demand information.

So, I separate one of the two effects by considering a middle treatment where I have firms being able to monitor what everyone else is doing, but they don't know -- but there's uncertainty about the demand schedule.

So, let me show you what subjects actually see to give you a better feel for what the design looks like. So, very briefly, subjects go through intensive training and when they get to the part where they choose their quantities, they see these matrices. The left-hand side is displayed permanently to them. We use callers to tell them to distinguish their pay-offs from the rival's pay-offs and also medium, high and low demand schedules. The probabilities are always displayed and this particular decision screen is for people who anticipate what demand schedule they're going to face. This would be the Rotemberg and Saloner design.

After they make their decision, we show them their profit by highlighting the entry in the cell that corresponded to what they chose and the other party chose. So, I notice that they have perfect monitoring here.

The middle treatment where they have uncertain demand information, they know that they're going to face one of these three. They make a choice and after they make the choice, we inform them of their pay-offs by showing them this matrix and you can see that they can also infer what the other person did. So, they have imperfect demand information, but they have perfect monitoring.

So, the key thing about our design is that in the third treatment where they have imperfect monitoring and imperfect demand information is that this would be the resemblance of the Green and Porter paper, is they wake up with a profit, but there is uncertainty about what the other person did. In this particular case, the subject chose to deviate. He woke up with a profit of

1250, but he really doesn’t know whether the other person chose to collude and the medium state happened or the other person chose to deviate and the high state happened. So, this is the way we include the imperfect monitoring feature in the design.

Let me just say that we went through extensive sessions by having several parameterizations. We had sessions in which we calibrated the parameters differently to create variation in the data and also to test the robustness of our results. So, to give you an example, for example, a parameterization in one should -- the parameters in one are calibrated so that the Rotemberg and Saloner prediction should appear whereas in parameterization two, it shouldn’t appear.

We did something similar with Green and Porter where we have the two different parameterizations giving you different predictions. In particular, in parameterization two, you should observe price war of a length of at least three periods. We do other stuff -- I will get to that at the end of the halftime -- that deals with risk aversion of subjects. We actually measure risk aversion and make sure that our results are robust to the presence of risk aversion.

So, very briefly, regarding the first result, we call that -- we would suspect that monitoring and removing demand information would reduce collusion and, here, we have a graph of -- on the Y axis is a measure of collusion and on the X axis are the periods the subjects played with. The dotted lines and the gray lines, I'll get to that in a minute. But the interesting thing here is that we should expect the blue line, which is the full information, perfect demand, full information treatment, which has perfect demand information and monitoring, we should expect it to have the highest collusion. But it actually turns out that it is the medium one.

In other words, when we go from perfect demand information, we remove that perfect demand information. We jump to the red line which means that collusion actually increases. But when we remove monitoring, which is going from the red line to the green line, it actually decreases.

So, just as a robustness check, we also find this in parameterization two. You can see that removing demand information, which is going from the blue line to the red line really doesn't change much collusion. But if we do remove monitoring, that does reduce collusion. Kind of a little bit counterintuitive. I should mention that if we actually work out the theoretical incentives of the medium treatment where we should expect collusion to decrease, on the left-hand side, I don't know, the equation is messed up. It should read a little bit different. But on the left-hand side now, the equation has an expectation because you don't know what demand schedule you're going to be facing, and it turns out in both parameterizations, the left-hand side is smaller than the right-hand side, and that's why people are colluding more than colluding less.

So, testing the two theories. Rotemberg and Saloner, remember that we have two parameterizations. In one we should expect it to occur and the dotted lines there are the periods when subjects are served a high demand shock. And as you can see, collusion drops importantly in the first parameterization and a little bit less importantly in parameterization two. So, graphically, at least we see some support for the predictions of Rotemberg and Saloner. This is confirmed in frequencies.

I also do an analysis of the subject level where I actually test how well the Rotemberg and Saloner predictions fair with other strategies, for example, the tit-for-tat strategy, which has been proven to have substantial predictive power, and other strategies like the finite punishment and the Green strategies. So, basically, I create an indicative variable which works like a robot. It takes a value of zero to one, depending on one of these strategies, and also one for Rotemberg and Saloner.

The random strategies just including our constant just basically says you choose to collude or deviate by flipping a coin. And the question that I ask of the data is how well do all of these strategies fit the data.

So, each column here has each of the strategies, and I'm going to focus on the load likelihood value that is at the bottom. And you can see that the Rotemberg and Saloner strategy and, also, the Green strategy perform best here, which I consider as a relatively not strong but somewhat supportive of Rotemberg and Saloner predictions.

Remember, these are tests on individual strategies. We can also do a test on outcomes, which would be the pair of strategies that firm chose in every period. What we find is that we create an indicator variable for each of those outcomes. In both letters, you see that the errors outcome would be both firms colluding when they're supposed to and zero otherwise. In parameterization one, that predictor does relatively well predicting 50 percent of the choices or the outcomes correctly. And in parameterization two, interestingly, they always collude at equilibrium, which is the one that we're supposed to be observing, predicts 71 percent of the outcomes.

Now, the Green and Porter theory, remember that is known as a theory that triggers price wars after a low demand shock. So, the gray lines that you see there represent periods when a low demand shock was observed, and we do see some drops in collusion or cooperation here when that happens. But we really don't see that pattern for which Green and Porter are known, finite regression to the Nash equilibrium and then back up.

So, we do a little bit more of individual analysis and we compare the Green and Porter strategies to other strategies. In this case, the strategies are going to be slightly different because since there is an imperfect signal, they really do not observe what other people are doing, we're going to be talking about thresholds. And the imperfect signal is the implicit price that they observe.

I consider two types of thresholds, one in which firms revert to the Nash equilibrium after they observe a load price and the two threshold strategy where they revert back up to the collusive level after they observe a sufficiently high price.

This is just a subset of the results, but here, I just want to point out that the Green and Porter equilibrium -- remember that it has many equilibria. But the one that predicts data best is the one that has N equal to infinity which is you deviate or start a price war of infinite period after price or the signal falls below a threshold level. And other threshold strategies explain data relatively well, too. But with the feature that firms stand, or in this case, subjects stand to start price wars of infinite length.

The test on outcomes, which is similar to the one that I did for Rotemberg and Saloner, tells us a little bit of the same story then the individual strategies. Green and Porter, with infinite price wars, predicts relatively well our data.

And just to wrap up, because I think I ran out of time, monitoring -- in this particular setting, monitoring appears to matter the most. So, when we remove monitoring, collusion drastically diminishes. This doesn't happen with demand information. When we remove demand information, collusion either stays the same or it increases, which is a little bit counterintuitive, but in line with the theoretical predictions for this particular setting. I think both theories have some support in the data. But the data kind of tells us that strategies tend to follow a Green strategy.

I think it's interesting to use experiments to sort out different predictions of different models, especially when we have models of such prominence. And in this particular case, we can rule out certain equilibria in favor of other. Of course, there are other potential applications, such as the merger guidelines, we know that one of the things that authorities have to look at are the potential factors affecting collusion after the merger. In this case, we know that it's not always the case that the stylized fact is true.

And I'm going to stop there because I've run out of time.

MR. SHAPIRO: The discussant is Joseph Harrington.

MR. HARRINGTON: Well, if I think about where I think the value of experiments in IO is particularly great is when it allows experimentalists where a theorist dare not tread. One such place is dealing with equilibrium selection. We know we're looking at collusion, we're looking at infinitely repeated setting. There's lots of equilibria and there's that question of, well, when is it that players, subjects, firms are going to be able to coordinate on something (inaudible) superior to a static or stage game Nash equilibrium, and when they can, what are the properties of that equilibrium.

That's something that theorists have something -- it's difficult to have really little to say about. And we know, certainly empirically, it's hard to say something.

So, I think this is a very well-motivated paper in that it's using experiments in an area where we do need some more insight and certainly theory has not been able to deliver that.

In terms of kind of the main take-aways to me, one -- I think it really comes down to this, that kind of pseudo-tastic collusion -- I say pseudo because -- and I'll come back to this point -- in that at the beginning of the experiments, there are some messages that are allowed to be conveyed between the subjects. But there are no messages over the course of the experiment, just preplay.

But, to me, the big take-away is that, in pseudo-tastic collusion, you know, you can collude with demand volatility. Actually, I think it's somewhat impressive that subjects were able to collude in the low to medium states, but not in the high demand states when it was appropriate for parameterization. But that with imperfect monitoring, it's a lot harder to collude, which leads me to kind of pose a question which I would put forth as a possible further treatment or another set of experiments, which posed the question of, when is -- if you ask the question of when is it that it's particularly valuable to explicitly collude, then I'm starting to think for these experiments that, well, maybe when imperfect monitoring is a real problem.

And we certainly know from a lot of different cartels that they have spent a lot of time and energy in terms of monitoring.  Lysine, vitamins, a whole bunch of them, they went to a big effort to engage in monitoring.

Now, here we -- I mean, these experiments can't deal with this question because you really need something where you have ongoing messages, but I think it's an important issue to address, which I'll come back to.

So, there's kind of two just points I want to raise about trying to kind of better understand the results.  One is to understand the fact that there's a declining frequency of collusion.  If you look at the two graphs, the one on the left is for the Rotemberg and Saloner full information treatment; the one on the right with the Green and Porter imperfect monitoring treatment. It's much more distinctive with the imperfect monitoring. But, generally, there's just decline in the frequency of collusion.  So, I'd really like to better understand to what extent that's an end game effect or to what extent that, for example, that they're using a grim punishment and we're just observing an accumulation of cartels that have collapsed.

A second point is about these messages. I mean, it kind of lets you know what was -- the types of things that were allowed for. So, at the start of the experiment, each of the players could choose a message from a limited set that might say, I'm going to play low every period, I'm going to play high every period. I will play low if only a few play low and so forth. And I can understand why you did that in order to be able to get more collusive equilibria. It would be helpful to report the results and how they related to the messages so that we can better understand the role of those messages. And, in particular, to what extent behavior was tied to whether those messages coincided.

I have kind of two comments about equilibria. One is dealing with this point here, this -- I'm quoting from the paper here, "contrary to conventional wisdom, removing demand information does not decrease, in some cases, it increases collusion." Just to kind of review on that, what he's contrasting is a treatment where it's -- a la Rotemberg and Saloner, you observe the demand realization prior to choosing your action. The alternative is you observe the demand information after the end of the period. So, monitoring is perfect in both cases.

Now, to me, I don’t know exactly where the conventional wisdom came. To me, it’s not surprising that you get more collusion with the ex post demand information because if subjects are -- certainly if they’re risk neutral, you’re just looking at something which is equivalent to determine a demand model. It’s just where you have expected demand instead of deterministic demand. But it would be equivalent.

So, then, if you contrasted deterministic demand with a Rotemberg and Saloner treatment, I would suspect deterministic demand would -- collusion would be zero. So, you’re finding that with a theory which is, I think, fully consistent with this.

The other issue concerns equilibrium and -- okay, so what was stated, delta equals -- it should be .75 here. That was kind of changed over the set of experiments. So, what he’s done is assume that there’s 30 periods for sure and then we’re going to start having random determination of -- at the end of the game.

Now, in characterizing what are the equilibria, he uses the discount factor of -- he uses that of .75, actually not .8. Now, the question is, is that appropriate given the fact that the discount factor is one over the first 30 periods of the game? Let's just focus on parameterization one. I think that is appropriate for the imperfect monitoring treatment because, in that case, when delta is .75, when you get to period 31 and your discount factor is .75, you cannot -- the Green and Porter strategy is not in equilibrium and thus, through kind of an unraveling argument, you're not going to be able to sustain that as an equilibrium for any of the first 30 periods as well. So, it's fine there.

Where it's more of a problem is with the full information treatment because, under the parameterization, what you show is when delta's .75, you can collude in the low and medium demand states, but not the high demand states. Now, that may be true when you get to period 31, but it isn't clear to me that it's true early on in the horizon when the discount factor is one.

Just in terms of future directions, one thing I think it would be interesting to do would be allow for a public correlation device. I mean, in essence, you have that with the full information treatment, which is through the demand. Now, the demand's also affecting the pay-offs but you can also use it as a public correlation device.

It would be interesting to have that in the imperfect monitoring treatment because one of the things that you find that's difficult is that once subjects stop colluding in imperfect monitoring, they have a hard time getting back to collusion. And you find the grim punishment is actually the best fit. If there's a public correlation device, that might allow them to get back to collusion.

The other thing I'll just mention is related to something I said at the beginning, which is kind of a broader point. I think a really important area for experimental work in relationship to IO is to get at this issue of explicit versus tacit collusion, specifically when is it particularly valuable that firms explicitly collude, engage in direct communications as opposed to tacitly collude. That's something which we have a very hard time providing any insight on theoretically.

Where I think it can be done experimentally is we look at a host of different environments. For example, here's two environments looked at by -- three environments looked at by Christian. And then to look at those under two treatments, one where they aren't any messages over the course of the experiment and one where there are messages, and you have to be somewhat specific about what kind of messages you're going to allow for.

But those types of experiments would be able to start getting at the question of when is it that we really think it's important to firms that they engage in direct communication? When does it have a lot of value? So, I'll stop there.

MR. SHAPIRO: Thank you. Again, there's a lot here. Are there some questions if you want to voice them at the microphone?

**(No response.)**

MR. SHAPIRO: All right. Well, please join me in thanking all the panelists.

**(Applause.)**

MR. SHAPIRO: I guess Chris will tell us what we do next.

**(Paper Session Three concluded.)**

**PANEL SESSION TWO: ECONOMICS OF PRIVACY AND INTERNET BEHAVIOR**

MR. ADAMS: So, what we're going to do is we're going to have a panel here. We have Susan, Amalia come up. Is Susan here? She just stepped out. Would somebody grab her and bring her back in?

The panel is going to be here. So, if people in the far reaches of the room want to move up so that they're closer to the panel and can actually hear and see what they say, then that might be of value.

So, we're going to have Amalia Miller from the University of Virginia; Pablo Chavez from Google; Sue Glueck from -- I'm not sure how you say your last name, you'll have to introduce yourself -- from Microsoft; and Alessandro -- and Alessandro will have to say his last name as well. He's from Carnegie Mellon. So, if we move up and we'll try to get this panel discussion on behavioral advertising and some of the privacy issues both on the Internet and more generally, get that going.

**(Brief pause in the proceedings.)**

MS. ATHEY: Welcome to the privacy panel. We have a very exciting set of people here today to talk about various applications of privacy regulation and policy. So, without further adieu, since we're kind of running let, we'll start with Sue Glueck from Microsoft

Corporation who's going to -- she's an attorney and she's going to be speaking about privacy issues in online advertising.

MS. GLUECK: This is always the challenging moment because I'm from Microsoft. If I make a mistake now, my credibility is destroyed.

So, what I'm going to do is try and set the stage so we can have a discussion about whether regulation really makes sense from an economic perspective for privacy on the Internet and, more specifically, for online advertising. So, I'm going to explain to you how it works and, hopefully, maybe even scare you a little bit. I know it's towards the end of the day, so I need to keep you awake.

We'll start by talking about the players. Advertisers, you know, those are the ads that you see when you visit Web sites, particularly when you go to visit free Web sites, like maybe you read your newspaper online so you don't actually have to pay for it. Well, it needs to be paid for some way.

So, the deal you're actually making is to see some advertising in exchange for free content. Oftentimes, that advertising is provided by a third-party ad network that provides you advertisements as you move around the Internet. Now, there are a number of them, but the same ad networks do see you as you visit various Web sites. So, I'm going to walk you through some scenarios so you can see how this works and what data gets collected.

The simplest scenario is when an advertiser just wants to buy some inventory or add impressions on a single site. So, the only data that the ad network server needs in that scenario is the IP address of the user, because, of course, that gets left behind every time you go anywhere on the Internet. So, the first party site that you visit, in this case, it's Kelly Blue Book, they get your IP address. But the ad server gets it as well. They see where you've visited, the KBB.com, and the time of your visit and they track what ad that you saw at that point. Pretty simple.

It gets more interesting when the advertiser is buying inventory on multiple sites. Let's just suppose that they only want to show a particular ad once to any particular customer. So, they want everyone to be net new to the ad. That means that the ad network server needs to track a cookie. They assign you a unique ID number. Now, they know you're user number 12345. So, they'll recognize you when you visit the next website. So, now you've seen this great ad from Chevy and they've recorded that you were on KBB.com and which ad was served and when.

Now, you're visiting MSN.com. They're not going to show you that ad again. They're going to show you a different ad. They know this from the cookie ID and from the other information about where you've been and where you've seen ads from that particular ad network. Over time, this accrues. So, the places you surf that are served by the same ad network wouldn't necessarily even need to be an ad network. If you see the same weather information on a website, that's coming from another server somewhere. That server knows that you visited three sites today and saw their weather on all three of those sites.

For targeting, sometimes what you want to do is target people who have specific interests. So, you're going to look for the users who have visited Web sites that reflect those interests. That's one way of determining what people are interested in. So, if you visit a lot of financial Web sites, then you're going to start seeing financial ads on the general purpose Web sites, because they know that you're interested in that.

When I started shopping for a car, first I went to Kelly Blue Book to look up the value of my old car. Then I started reading about cars online on various Web sites and looking at, you know, how much should a car cost, and I started to notice that wherever I went, I was seeing ads for cars. The interesting thing is fairly quickly that went away because within a certain number of weeks, most people make their car purchase decisions. So, they don't show you those advertisements anymore because you're probably sick of looking at cars. You're done with that.

So, you are getting some value there. You're seeing advertisements that are more relevant to your interests and what you're looking for. You know, if you're shopping for a camera, those decisions are usually made very quickly, I think within a week or so. So, you couldn't see camera ads or cell phone ads for too terribly long.

I think this is something that people generally don't -- none of my friends knew about this until I explained it to them. Let's say you go to Walmart.com and you put some stuff in your shopping cart. Walmart advertises on a lot of Web sites. So, what they do is instrument their Web sites as the publisher. They instrument their Web sites so a Web beacon indicates that, hey, this is a Walmart shopper and let's say you abandoned your shopping cart and, in this instance, it was full of maternity clothes, that way the ad network can know that about you. You're someone who shops for maternity clothes and someone who shops at Walmart. So, then when you go to some other Web site, you're going to see an ad for Walmart to try to bring you back to finish that purchase.

Meanwhile, this data continues to accumulate. Where you go -- and don't worry, I'll tell you how to make it stop in a couple of minutes.

**(Laughter.)**

MS. GLUECK: You know, where you go, what you're interested in. This may even include some registration data that you gave when you were registering on a Web site. So, for example, if you register for a Windows Live ID because you have a hotmail account or you use MSN Messenger or something like that, we collect a little bit of demographic data about you. I think gender and country and something else that's not popping into my head. Pardon? Age, thank you. Although you don't always get asked age. It depends on where you are in the process. But a lot of times age so that we can do the COPPA screening if a child hits something that would be -- where COPPA is relevant.

So, we're very careful to take that demographic data we have about you and we associate it with a number, but we never associate it back. So, your name, your hotmail email address, any other email address you provided us, that's never associated with the advertising data. I can't tell you how everyone does it, but I think we're somewhat unique in that system, and I'll tell you a little bit more about that later because it helps to make our opt-out of targeting advertising cookie a little bit more effective.

So, as you think about this, we saw a lot of different types of data that was accumulating in the profit and associated with a cookie ID. Over time, you know, if you know enough about a person, they may become more and more identifiable. Search queries, that's something else that could be in a profile, and Pablo will be talking about that later and how that works.

So, if you think about it, well, gee, is this really a big deal? A lot of this data is pretty innocuous. Well, if all third-party ad networks were sort of created equally and, you know, this ad network had a little bit of data about you and that one did, it's probably no big deal. But as the ad networks themselves have more and more market share, then they know more and more about you. They may not know your name, but they certainly know a lot about your habits.

DoubleClick, I believe, is the largest in the market. I think they have about 70 percent. They're owned by Google and Google also has AdSense, which is another advertising serving mechanism. So, Microsoft has a sizable network as well. It's called Atlas. So, over time, you know, companies are accumulating more and more data about users, which raises, of course, some privacy concerns.

I'm always asking, as I visit Web sites, well, gee, what are they doing with my data? How long do they keep it? Do they anonymize it? And if they anonymize it, what method do they use to anonymize it? At the end of -- we retain search data currently and advertising data for up to 18 months. At that point, the IP addresses are all completely wiped out. The cookie IDs, any cross-session identifiers are completely wiped out. So, we don't know -- we know somebody searched for maybe this address, but we don't know who it was or what other things they searched for at the end of that period. Different companies handle this in, of course, different ways.

I really started thinking a lot about anonymization when back in 2006, AOL had a data breech where a well-intentioned researcher posted 650,000 users' search data over a -- from three months of search data from that many users, which turned out to be about 20 million searches. And they took out the user names so that, okay, I've anonymized the data. Well, there were still cookie IDs. So, if I had done 200 searches during that three months on AOL, then you would know that User 12345, what all I had searched for.

It turns out a disturbing number of people in the United States are thinking about killing their spouses from looking at the searches. I am completely serious.

**(Laughter.)**

MS. GLUECK: Or maybe they're writing murder mysteries. If you're a glass half full kind of person, maybe there's a more positive spin you can put on that. Two people were identified fairly easily by the press. So, here we have this anonymized data. People were identifiable. There were other examples. After the session, I can regale you with them all night long. But it really makes you wonder how anonymous is anonymous and how much do you know about what people are doing.

Well, the industry is self-regulated. The Network Advertising Initiative in 2002, but thanks to some concerns about DoubleClick, was formed. Atlas, our ad network, was one of the founding members, I believe. And there are rules for how you do these kinds of things, what kind of consent you need to get from people, for example, if you're going to advertise to them based on sensitive information like health, that you get consent for doing that, because that's a little -- you know, just because a friend tells you they're sick, you look up a medical condition and, suddenly, you start seeing these ads, that would be a little creepy.

So, the FTC, this year, proposed guidelines, which is great because, you know, after six years, that seems very timely to take a look and say, does the industry still look the same, are the players still really the same, you know, what about inspection. There are new questions today. And, so, they've proposed guidelines for self-regulation. Three states have proposed legislation, New York, Connecticut and Massachusetts. And I have to say the legislation is pretty well-crafted. Clearly, they had an understanding of how the business works, how the technology works and what data we're talking about.

In Europe, we think they'll probably turn to this next, they've been -- the data protection regulators have been very focused on search, and so, it's sort of the logical next thing for them to start looking at.

So, I should have disclosed at the beginning, I'm not an economist, I'm a lawyer. And I apologize, I have no charts or graphs for you. But, you know, I start to wonder how much users understand about their privacy. I mean, how much did everyone here, did you know that it’s possible that you were being tracked? Well, the good news is, there are things you can do about it.

If you go to the NAI website, you can go and opt out of all the member ad networks at one time. So, you can see this is just the very beginning on the screen shot, of the list of networks you can opt out of being targeted. Your information may still be collected, but they don’t use it to actually target ads to you.

Internet Explorer lets you control cookies. It’s a little hard to live in the world and use the Internet without accepting any cookies. So, I think that’s a little harsh perhaps.

Reading privacy statements, always helpful. We’re big fans of the layered notice approached where you provide meaningful detail at the top layer. In this instance, you can go directly to our display of advertising section from the top level and learn how to opt out of getting advertising from Microsoft, but targeted advertising from Microsoft. You’ll still see sort of generic ads coming up, but the ones that are just for you because you are a woman who -- if you’re me, you’re a woman who’s 46 years old and that kind of advertising stops.

The neat thing about our opt-out, if you check the top box, then it’s just like everyone else’s and if you're a little paranoid about your privacy and you delete your cookies regularly, then your opt-out cookie goes away, too, and it's very sad, you start to get targeted advertising again.

If you check that second box and you have a Windows Live ID, then every computer you log into using your Windows Live ID, the targeting will stop. If you get rid of all your cookies, all you need to do is log back in, which you might be doing to check your hotmail or, you know, just -- I don't think people log back in just to set the cookie. I think they do it as a natural part of doing other things. But it sets the cookie all over again. I believe this is unique in the industry.

Internet Explorer 8, we realized the beta version of this some months ago, and it's got this new thing called in-private blocking, that lets people block third-party content. So, it's not just ads. It's a map, a stock ticker, the weather, things that aren't coming from the website you think you're visiting, you can block or you can choose to always allow, which, you know, I personally think it's a good deal to get to see free content in exchange for advertising. So, I'm not blocking anything at the moment. But it's nice to know that I could because I'm using the software.

The difficult part in working with the --

because I support this product -- in working with the team was how do you explain to end users what they're blocking and what they're allowing when they don't have the faintest idea, some of them, that third-party -- a lot of first-party Web sites serve third-party content?

So, again, we're sort of back to that question of, you know, regulation, does regulation make sense in this area? And it may. Of course, if it took into account how the industry works, how the technology works and allowed for the ability to use data and innovate using data, you know, good regulation, might make sense to help protect consumers. Bad regulation, I think, would just hurt the industry.

MS. ATHEY: So, why don't we keep moving through the panel in the interest of time and then we'll come back and have questions for everyone.

MS. MILLER: Mentioning regulation is a good segue for what I'm going to talk about. I'm just going to spend a few minutes telling you about some of the results from some research that I've done with Catherine Tucker at MIT, looking at the effects of privacy regulation at the state level on the diffusion of a particular form of health information technology, in particular, electronic medical records.

So, there's sort of this question about good regulation, bad regulation. In some sense, we want to ask, looking at the regulations that have actually happened and that exist at the state level, about half the states in the U.S. currently have some form of additional requirements that restrict the ability of healthcare providers to share information that they have among each other without express consent from patients. So, if you have some private medical information that your hospital knows about you, half the states have some extra standards above the federal minimum standards protecting the privacy of that information.

We want to look at what happens or what’s the effect of these regulations on the diffusion of electronic medical records. Electronic medical records are this technology that basically allow you to use computer systems instead of paper records to keep track of, so to store medical information, to retrieve it within a hospital. But also one of the key benefits from this technology, one of the key promises is the ability to exchange information across healthcare providers faster and more cheaply.

So, when you think about what privacy rules might do to the diffusion or to the benefits from the point of view of a hospital of switching over to electronic records, you can imagine sort of two possible scenarios. It's possible on the one hand that regulation is inhibiting diffusion. The way that that would happen is that the regulation puts a cost -- every time you want to share information or share information, privacy protection can make that more costly and more difficult.

So, if you're thinking about adopting, you think about the network benefits, the benefits that you'd have from other local providers or even more distant providers who also have electronic records, that benefit of being able to exchange information about patients more easily is going to be reduced when you have to overcome an institutional or regulatory burden. So, that regulatory burden is going to replace the burden you had in terms of the physical challenge of exchanging this medical information and might make hospitals less likely to adopt medical records.

On the other hand, it could be that patients are very concerned about privacy. They might be concerned for reasons that Sue mentioned. They might be concerned about having their identity stolen. Medical identity theft is a new phenomenon. They might be concerned about having their neighbors or coworkers find out about health problems that they have that might be embarrassing. So, they might not want to go to a hospital that uses electronic records. They may not want to have their information stored in a way that could be easily accessible or that may be more vulnerable to exposure.

In that case, it could be that when a state comes in and says we're going to protect your privacy, we're going to put some strong regulation in place, that might make consumers feel more comfortable and more safe with electronic records and that might promote adoption. So, there's sort of this potential cost or benefit and we don't really know what the net effect is going to be. And what we do in our paper is we try to empirically assess which is, in fact, the case.

Just a bit of background in terms of why we're interested in electronic records in particular, this has been something that politicians in the U.S. have been talking about for a very long time. Healthcare information technology and electronic records, part of that has been lauded by politicians across the spectrum as this great technology, this great innovation that's going to both reduce costs and improve outcomes.

So, we have a quote from Newt Gingrich and Hillary Clinton both agreeing that healthcare IT is great. You heard the last presidential debate, Barack Obama and John McCain both also support adoption of healthcare IT. So, this is a technology that's been around for decades since the '70s and adoption in the U.S. is still pretty low. Under 50 percent of hospitals have switched over to electronic records. This is a concern, and so, we want to know if privacy is having a role in terms of either -- possibly slowing that adoption.

This is just more stuff about why -- so, this Bush Administration had a target of national EMR adoption by 2014. People are skeptical about whether or not that will happen. The Federal Government's been very concerned about privacy. Consumers have expressed concerns about privacy for electronic records. The government attention, to date, has been to try to figure out how to make privacy standards tough enough. So, there's a $17.3 million study that was trying to assess how can we ensure privacy.

There's a lot of media attention that talks about kind of what happens when privacy fails and when -- especially information about celebrities is disclosed. So, George Clooney was in a motorcycle accident and everybody heard about it. Britney Spears went to rehab and we knew about it. But there's not a lot of discussion about what the potential costs are from imposing strong regulation.

And the particular cost that we're thinking about is this trade-off where strong regulation that protects privacy might be blocking these network effects from sharing medical information. As far as we know, nobody's looked at the other side of this and that's what we're trying to contribute.

There is some anecdotal evidence, other than our study, that has -- where vendors have said that strong privacy laws can be a challenge and, also, medical providers have sometimes said that complying with complicated state regulations, in terms of protecting privacy, have led to a particular regional effort to combine and share health information. In Southern California, it actually fell apart after several years and a lot of money went into trying to create it and the regional initiative fell apart and the large -- and the participants mainly blamed the challenge of trying to comply with the strict California state privacy rules.

So, in terms of everything that I'm going to talk about now, our results -- our empirical study is looking at this particular technology, healthcare IT. We think that some of these trade-offs in terms of privacy and technology adoption might have some implications for other types of technologies where there are network benefits that have to do with sharing information.

So, the data that we have, we basically need data on two components. One is the adoption decision. And we get that from the HIMSS Dorenfest database. So, we have data on adoption through 2004. And we can match that with some information from the American Hospital Association to learn a bit more about the hospitals.

The period that we look at is from the 1990s through the end of 2004, into 2005. And you can see that that's the period when most hospitals in the U.S. are adopting. This is just a histogram of the number of new hospitals adopting and that 1992 bar is the total number of hospitals that adopted in '92 or earlier. So, really, this is the interesting time period to be looking at to study adoption of EMR.

Then we need to combine that with some data on privacy laws, which we get from a group of Georgetown University called the Health Privacy Project, whose function is really to understand privacy laws and, also, to advocate for stronger privacy protection and consumer protection. So, we get laws from them. We have a panel of laws. This is just the cross-section in 2000, so you can see that there's a lot of variation. About half the country, half the states have a law; half of them don't. There's no obvious red state, blue state configuration. All kinds of different states in all different regions, some of them have and some of them don't have privacy laws.

So, the variation that we use is a variation across states and over time in privacy laws. And without showing you any tables, even though I am an economist, no tables and no figures, I'll just tell you the results. We find that states that have strong privacy protection, this leads to a reduction in EMR adoption. And by 2005, it leads to a reduction by about 24 percent. So, it's pretty substantial cost in terms of slowing technology adoption.

When we think about what the mechanism is or the channel for why are these laws inhibiting adoption, we find evidence that it is, in fact, these network benefits. So, when we look at the correlation in the adoption decision or the responsiveness of one hospital's adoption to other local hospitals' adoptions, in states that don't have privacy laws, we find substantial positive response. When other hospitals in your area adopt, you're more likely to adopt the same technology. In states that do have privacy laws, that effect goes away. So, there's no responsiveness to other hospitals' decisions and we think that that's suggesting that these network effects have been substantially diminished.

We also look at the effects, not just on adoption at all, but on the types of technology that hospitals adopt.  So, there are a lot of different vendors that make electronic medical record systems. Some of them are designed to be interoperable and compatible with systems made by other vendors.  Some of them are meant to be closed loop systems that don't play very well with others.

We find that privacy laws not only inhibit adoption, but they also lead hospitals to adopt systems that are less compatible, less compatible generally and less compatible with the systems adopted by other hospitals in their area.  So, when the law is coming in and telling you you can't talk, you're less interested investing in a system that is able to talk well with the other systems in the area.  And those are the results.

We have another paper where we kind of make an argument that some of this adoption is actually leading to adverse health outcomes where we find that adoption of electronic medical records actually reduces infant mortality rates.  So, we do find some evidence that these electronic medical record systems are improving health outcomes.  So, that's a reason why we're very motivated to think about this adoption and why this is a cost that we should weigh against the benefits of potential privacy regulation.

MS. ATHEY:  And, now, it's going to take

Alessandro just a minute to get his presentation loaded up. So, now we’ve heard about two specific places where information can be used to provide a lot of benefits and privacy laws can potentially get in the way of that. Now we’re going to hear a little bit about how consumers think about privacy and whether they actually understand what they’re getting into. Then, finally, we’ll hear from Google about some of the -- as they kind of pull some of these ideas together.

MR. ACQUISTI: So, my research focus, I call it the economics of privacy and the behavior of economics of privacy. It’s a study of the trade-offs associated with the protection and the revelation of personal information and the study of how individuals make decisions about those trade-offs, decisions that sometimes may sound contradictory or even damaging.

In fact, let me start with one shot from the Daily Mail 2007 about a Facebook group called 30 Reasons Why a Girl Should Call It a Night. So, Facebook is an online social network on which 90 percent of our students are, and no longer only students, also people after college are on Facebook. What is interesting is that not only people reveal personal information such as birthdate and sexual and political preferences, but in some cases, they also reveal information which could be embarrassing or damaging.

In this particular group, ladies post photos of themselves in various states of being passed out or drunk and sick from drinking and so forth. It's not angry, angry ex-boyfriends posting this information. It's the person herself posting this information.

So, why? We could conclude that the Facebook generation has no sense of privacy whatsoever, no need for privacy.

**(Laughter.)**

MR. ACQUISTI: But, in fact, that's not really necessarily the case because pretty much at the same time when the article was published in the Daily Mail, also in 2007, this other article came out about the Beacon. You may remember Facebook started pushing for this advertising program called Beacon, which would gather more information about Facebook users and spread it around to other Facebook users. Sure enough, around 700,000, apparently, Facebook users reacted violently against the Beacon and forced Facebook management to go back in their plans of pushing for these Beacon strategies.

So, what we have at the same time, we have some need for publicity, even bad publicity, and the need for privacy. They seem to be contradictory needs, but, in fact, they exist in each or every one of us and they only show control. Privacy is often defined as a control on what is public and what is private. But, for myself, my background is economics. I can really control this initial signaling.

The lady who's publishing photos of herself passed out and drunken is signaling information to a certain peer group, right? To a peer group in which being passed out means that you can party hard, you're a fun person and so forth.

Well, the problem is that you don't have control of the information once it's put out there. You no longer can know who else will see that information. Maybe your parents, maybe your future employer. Maybe that information is cached somewhere, and 20 years later, when you are going for Supreme Court candidate or maybe a Vice Presidential candidate, the photo pops up again.

So, there are costs and benefits in revealing, as well as protecting information. That's what the economics apprise us about. It's not new. Chicago School economists were the first dealing in this area -- with this area, Stigler and Posner, but yet they -- Chicago approach which was privacy sometimes creates inefficiency in the marketplace because it reduces information. Varian, Noam, Laudon in the mid-'90s introduced more IT expertise into the economics of privacy. And then, more recently, after 2000, a number of people, (inaudible) Taylor at Duke, (inaudible) Pavan and myself (inaudible) and Calzolari (inaudible) Berkeley, started working with microeconomic models of privacy, especially privacy (inaudible).

But as we were aptly modeling a way, we were in a two-peer model, we have a high and low consumer buying goods and (inaudible) tracking them and trying to learn personal information to (inaudible) them, we learned also something surprising, which although people were claiming that privacy is important, you need more privacy, if I had more privacy, it would show up more aligned. If I have more privacy, it would go more aligned. In fact, behavior did not reflect those attitudes.

So, we started discovering, as we were doing these models, that reality was telling us that people want privacy, but they don’t want to do anything about it. They rarely pay for it. In, fact they can be convinced very easily to clear away lots of personal information for a small reward. This was shown by (inaudible) Spiekermann in Germany.

More recent studies that we did on online social networks showed a clear dichotomy between what people say they want to keep private and what they do on online social networks. Leading many people to say -- to ask, so, do people really care for privacy? Is it something that is important? And if people care about it and they don't do much about it, should we be (inaudible) enough to take their own protection, so the FTC or policymakers or businesses should protect what consumers are not protecting themselves. So, how can we answer that question?

Well, this trying to answer the question led me to engaging into behavioral economics of privacy. Because I realized that taking a fully rational approach here would not address all the issues. The fully rational approach would be the one in which Johnny MySpace is thinking whether he should reveal or not certain particular sexual kinks on MySpace. And he's thinking that, well, if I do so, I will find somebody who has similar kinks as I do. That's good. But maybe my future employer will find out these kinks and will not employ me or maybe my parents will see it. Therefore, Johnny will solve a complicated question.

**(Laughter.)**

MR. ACQUISTI: Over time, this kind of (inaudible) cost him benefits. We don't do that, really. We use (inaudible).

This is a similar example. This is a real flyer which was handed out at the San Francisco pier when I was a student there, so around 2002, 2003. And it asks people personal information, to fill out this form indicating age, mental status, occupation, income, credit card, even address. In exchange for participating in a lottery with the odds of one out of 700,000, winning around $25,000. So, there are economists here, so you can easily do the expected value. It's basically a few cents. Not even worth the actual opportunity cost of spending time filling this form out.

But the problem is when people see this form, how can we make it a truly rational decision about what is the best approach. Should we fill it out or not? What is the difficulty of making this decision? Difficult framing, that even if we care about privacy, we care about privacy in general. In the specific, well, yeah, with specific benefits and specific costs, we may say we want to protect ourselves, but then we don't want to spend time maybe changing the privacy settings or the cookie settings on the browser. They are there and they cost only 10 seconds to change, but those 10 seconds are too much.

Incomplete information, so things that we don't even know that the problems are there and we don't know that the solutions, such as changing the cookie settings, are available.

Boundary rationality, Facebook has a very granular -- gives very granular control to users to decide what to reveal to whom and when. It's almost too granular, in the sense that behavioral economists know the paradox of choice. That sometimes when you give an incredible amount of different sections to choose from to users, the final decision could be sub-optimum.

And, finally, even if we had complete information (inaudible) there are all these psychological behavior biases that experimental behavior economists have studied over many years that seem to all apply to privacy decision-making. Indeed, in my final minutes, I will show you one particular study of the many we're doing with George Lowenstein and Leslie John (inaudible), one particular study in which we basically take one idea from behavioral economics and we apply it straight to privacy.

So, the first two, I will not discuss because I don't have time, but they are about how you can frame differently a certain survey or certain questions and impact the propensity of people to reveal personal information, as well as when you give assurance of privacy to people paradoxically, they start revealing less. When you tell people their privacy will be protected, they start lying and revealing less probably because they become frightened and (inaudible) about the sensitive (inaudible).

The study we focus on is on the effect of framing of privacy (inaudible). So, this is the story. So, basically, it’s a traditional endowment study, only that we tried to translate a downward study into privacy (inaudible). In privacy, you have willingness to pay and willingness to offset. You are willing to offset money for your privacy when you’re searching information on the Internet because you are using a service, but you are revealing information about yourself. Your IP address, your interests and so forth. You are exchanging something for your data.

But sometimes the results (inaudible) protect when, for instance, you decide to go and delete your cookies. There is an intangible cost, the time you spent to clean up your system, that’s the cost that you are offsetting to engaging.

So, the framework we wanted to use to study this problem was -- and this was an experiment we did in the field. We stopped people in the mall and then randomly assigned these people to different groups. In one group, we told people, hey, would you like to participate in a study, and if you participate in this study, you will get this gift card that you can use where any debit or credit card is accepted. By the way, this gift card is worth $10 and it's anonymous. Whatever you purchase with this card, we will never know, nobody will ever know.

Then there was a study which was completely unrelated to us. We didn't care about the actual study. And then something (inaudible).

To another group, experimental group, instead we said, hey, would you like to get a gift card to participate in this study. The gift card is valued at $12 and your name will be recorded and the transactions you make with this card will be known. People got the card. Did another study. Again, this study was completely unrelated to us. Eventually, each group was offered to swap cards.

The first group who was given the $10 anonymous card was told, hey, by the way, thanks for participating in the study. Would you like to swap your $10 anonymous card for a $12 identified card? In other words, would you like to get two more dollars to give away your data. And the second group instead was told, hey, by the way, look, you accepted this card, $12 identified. Would you like to swap it for a $10 card which is anonymous? In other words, this group was told, would you like to give back -- give away $2 to protect your data?

Remember that rationally the two groups exactly the same (inaudible). Do I want a $10 card and privacy or do I want the $12 card and no privacy? So, in the group which started from the $10 card, 52 percent decided to remain with the $10 card. Kind of status quo (inaudible). They decided to remain and 48 percent decided to switch. In the group which started from the $12 card, only 9.7 percent decided to switch to the $10 card. So, this difference you can guess is strongly statistically significant and is not only due to status quo bias. There is something more. There is the fact that when you frame differently a privacy money problem, people will give very different answers.

To conclude, this is precisely the point I'm trying to make with this research. That privacy evaluations and privacy decisions can be manipulated, can be affected, and therefore, this is something that should be considered because we cannot, I believe, rely on the (inaudible) preference in what people do and conclude that people don't care about privacy. Maybe it's more complicated than that and, therefore, often this could also be considered at the policy-making level because definitely companies are considering this because of the ability of getting information easily from people.

I should stop here and let Pablo talk next.

MR. CHAVEZ: So, next time don't show my sister's profile.

MR. ACQUISTI: I didn't know your sister was from the UK.

**(Laughter.)**

MR. CHAVEZ: Hi, good afternoon. So, first of all, thank you very much to the FTC and also to Northwestern University for giving Google the opportunity to participate on the panel. Thanks for just great colleagues and great presentations.

When I first talked to Chris about participating on this panel, really the notion was to look at privacy in a broader context and to look at the variables beyond privacy that a company would look at in order to determine what privacy protections it's going to provide to its users. So, from the get-go, I just want to say that we take privacy very, very seriously. My goal here, though, is to give you a sense of other variables that any given organization, by the way, including government, might want to take into account in determining what rules of the road to apply in any given situation.

The second thing that I want to accomplish is just to give you a snapshot, a window into Google's approach to privacy, which is just, actually in many ways, very similar to our approach to product development. And, finally, hopefully, this serves as a framework, not the framework, but a framework to look at a potential privacy self-regulation regulation and legislation that may be coming up. Just one way to look at it.

Lastly, I guess a couple of things I'm going to present might actually serve as good raw data for you guys maybe for next projects. So, hopefully, it's helpful in that sense.

So, first and foremost, Google, like companies like Microsoft and others, is very, very focused on innovating and we innovate through iteration, so repetition and experimentation. That is also very much our approach on a good number of privacy issues. When I'm talking about privacy, really I'm talking about the collection, use and retention of users' personal information, so maybe narrowly defined.

In this case, what I'd like to talk about a little bit is retention of our search logs. So, first, I'll give you an explanation of search logs themselves; secondly, I'll talk about our retention policies for those search logs; and finally, I'll give you a sense of kind of the other factors that needed to be taken into consideration as we determined kind of the retention policies.

So, very simply, if you were to go to Google.com right now and type in a search, Federal Trade Commission, for example, we would collect standard log data, much like what Sue talked about and others have talked about. This is the URL, including the search query, in this case Federal Trade Commission, the IP address associated with the device from which the query originates, the time and date of the query, the operating system of the device, the browser type, so I8, Chrome, Opera, Firefox or whatever you happen to be using, and a cookie ID.

I'll note, too, that this is a situation where an individual has not signed in to Google. So, this is an unauthenticated individual who's conducting a search on Google.

So, what we've decided over time is to reduce our retention or anonymization period for search logs for server logs generally to nine months. What I mean by that is that essentially, as you can see in this graphic, which by the way, we presented our privacy policy just to give users a sense of what exactly we're talking about. We delete the last octet of an IP address associated with search query after nine months.

Now, back to the iteration point, Google started as a company that retained this data indefinitely. So, the company’s 10 years old now from soup to nuts, from the moment of incorporation until now. For about eight years of that time, Google was retaining these logs indefinitely. Now, various stakeholders, policymakers, regulators and others expressed concern that, in particular, IP addresses were considered to be personally identifiable information. Google, by the way, takes a little bit more of a nuanced view. We believe that IP addresses are PII, personally identifiable information, in some circumstances and not in other circumstances. I’m certainly happy to answer questions about that.

So, we went from indefinite to about 20 months ago a policy of anonymizing after 18 to 24 months. Shortly after that, we went to a policy of 18 months. And in September of this year, we went to a policy of nine months. Essentially, it’s taken time to really look at the effects of these shorter retention periods on various services that we provide, on security issues that we have, on integrity issues that we have with the system.

So, just to talk about a couple things, some of you may be familiar in 2004, our system was attacked by the Santy worm. It was a pretty violent thing. Essentially it was very, very helpful to have IP addresses in order to combat that worm to understand essentially what device, what server had attacked Google. Now, this had nothing to do with the geography, it had nothing to do with who was operating the device. Rather, it was just making sure that whatever a device was sending us the bad request needed to be blocked. So, that's one example of a security application.

Many of you know that the companies like Google experience click fraud. One way that we fight back against click fraud is by looking at IP addresses where, for example, there might be a pattern of behavior associated with an IP address that would indicate that it's a fraudulent click. As some of you know, advertisers pay when you actually click on the ad.

So, these are some of the examples that really kind of drive this and a sense of balance of factors that we consider as we look at data retention policy. Again, one piece of the bigger privacy issue.

Let me talk about the up-side a little bit of data retention, and in this case, again, I'm talking about IP addresses. So, if you've used Google search, you've probably noticed that there is a spellchecker function associated. So, there's a feature that suggests

-- respectfully suggests the proper spelling for a query. That is developed -- in part, one of the signals that you're looking at is, for example, in a particular session you could have somebody spell apple with one P and then spell apple with two Ps and we know that the device from which the query originated probably misspelled it the first time and then spelled it correctly the second time. That is then inputted as essentially the correct spelling for A-P-L-E and it's suggested in future instances of that misspelling. So, it's a great improvement.

To me, though, what's really fascinating and something that I use and I would encourage you guys to take a look at is Google Trends. So, I've actually been looking, just doing a little bit of a study on electoral results. Essentially Google Trends is a tool that allows you to look back at some period of time and pick out specific search terms -- specific search terms made from particular geographics at particular times and figure out patterns and potentially connect those search terms at those particular times from those particular geographies with something that's happening in the real world.

So, what I looked at, I actually looked at the Iowa caucuses on the Democratic side in January of 2008. What I wanted to see was whether there was any correlation between the search queries relating to the three top candidates, in this case, Senators Obama, Clinton and Edwards. So, whether there’s any correlation between those queries and actual results in the caucuses.

So, as you can see, what I looked at is United States, subregion Iowa, January 2008. The caucus itself was around January 4th. So, a couple of dash marks to -- well, to my right of January 6th. And really, really interesting. You see that Obama -- so, Obama won 38, Clinton got about 30 and Edwards got about 30. So, you see Obama trended upwards significantly prior to the caucuses. And really interesting, you see that Edwards and Clinton were basically tied just like in the electoral results.

So, this is just one application of Google Trends, which uses, again, IP addresses to a certain geography. You can see that we’re geo-locating to the level of city here in a very, very useful way to consumers, in a very useful way to researchers. You can imagine other applications, for example, looking at health trends throughout the world, whether search queries can actually tell you something about something that has happened historically in the health area or currently in the health area.

One other potential application is the effect on the economy. Is there potentially a correlation between searches for cars in, say, the eastern United States and an up-tick in car purchases? Could that actually tell us about an economic recovery? So, just really, really fascinating data, again, tied to IP addresses and IP addresses have been kind of the focus of regulators and legislators as a potential area where companies should be obligated to delete IP addresses after some period of time.

Again, without passing judgment on whether that's right or wrong, there are definitely implications to such a policy.

I've kind of described the issue as kind of an either/or situation. You know, the truth of the matter is that when you're innovating around privacy and you're iterating and experimenting, oftentimes, it's actually a both-and proposition for consumers.

So, Sue had mentioned before our acquisition of DoubleClick. DoubleClick is an ad-serving company. So, it's one of the companies that's providing the display ads, the banner ads that you see on, for example, the New York Times or the Washington Post. One area where we thought that would be interesting to look at privacy innovation -- and we submitted this actually as a comment to the Federal Trade Commission's draft privacy principles -- was the motion that maybe we could actually provide additional transparency and choice to consumers when they receive one of these ads from a third party.

So, if you were to click on the ad itself up top, the landing page would be for a product or a service offered by that company. But if you were to click on the links on the gray strip at the bottom, that could potentially take you to a privacy policy, explaining what data is collected and how it's being collected, or potentially give you the opportunity to opt out of data collection, or give you the opportunity to actually comment on the ad and say whether you liked it, whether you ever want to receive anything like that again.

So, again, this kind of experimentation, you know, we would hope would be encouraged. This would be, from our perspective, kind of a great tool to adopt for industry at large. But this kind of an idea doesn't come up without kind of the opportunity to really experiment around privacy.

One last example of experimentation around privacy, Sue had talked about privacy policies and layered policies. We agree that those are great. One area where we're experimenting is really kind of expanding beyond the notion of a privacy policy and really kind of talking about privacy center where we associate kind of privacy practices and principles with a particular product. That's a lot of Ps.

So, in this case, we have a privacy center around advertising and privacy that talks about our privacy principles relating to advertising, that also explains products so people understand exactly what's going on, that allows them to opt out of data collection if they so wish.

Then, finally, and this, to me, is the coolest thing, we've been working very hard on a series of privacy videos, on YouTube videos. We actually have a privacy channel on YouTube and we also feature the videos in our privacy centers. That is a great -- that's actually an experiment that's turned out really well, because not only are we providing kind of five-minute, ten-minute snippets, plain English explanations of privacy, but also we're opening up these videos to comments.

So, no longer are we talking about privacy policies, specifically where it's a one-way conversation, we're just pushing out information. But, rather, we're collecting comments from our users, hearing from them about what they like, what they don't like, what they're comfortable with, what they're not comfortable with, and really kind of engaging in a dialogue with consumers.

So, again, happy to answer any questions. Thank you very much for your attention.

**(Applause.)**

MS. ATHEY: So, since this is a panel discussion, I'd like to sort of give a chance for each of the panelists to answer kind of a general question and then we'll throw it to the audience. I know everybody's kind of anxious to leave, so we'll try to move along quickly.

But I think that one thing that didn't come out completely clearly in some of the discussions is just how -- some of these options for privacy opt-out where the choice you were given was either you received the ads and you get these targeted ads or you opt out and you don't. But I think going forward, what the regulators need to be thinking about is how the regulations that they put into place are going to affect competition in the marketplace, say what's going to drive Microsoft to create an innovative product that makes it easier for you to manage your cookies and will that have sort of follow-on implications? Are we going to start seeing ad networks providing you with coupons or rebates to help induce you to participate? So, what's going to lead some of these benefits to flow back to the consumers.

So, I guess the general question I'd like to get people to react to is what kind of regulation do you think will both allow the innovation and the potential benefits of using information to go forward and also will lead to benefits flowing back to consumers?

So, maybe everybody can go down and react.

MS. GLUECK: Well, now I wish I had put in the slide about our cashback program. If you click on the ad, you can get some cash back. But I think it's worth looking at an EU style national privacy law. Companies -- you know, a lot of companies are already living up to those obligations because they're not just U.S. companies. They do business all over the world. So, the additional compliance cost for companies are likely not to be significant.

I think that really the important thing, regardless of what happens, to the extent there is regulation, that it's done very thoughtfully and carefully because I was fascinated by your work and the effects you saw, unwittingly, you know, and it -- I actually have a lot of questions for you about is there such a thing as a beneficial state privacy health law that could actually promote data sharing and, at the same time, not reveal when celebrities go into rehab.

MS. MILLER: So, what I can say from our research on privacy laws relating to healthcare, there is substantial variation in the state privacy laws. We actually don't know or we've sort of tried a little bit to see if we can identify dimensions that were more helpful or less harmful, and we didn't in terms of variation between the state laws. That's something that we're definitely interested in for future work.

But, right now, what I can say is that there is a federal law, there's a privacy rule as part of HIPAA that went into effect in 2003 and we don't find any detectable effects of the HIPAA regulation coming as a play. So, everything we're finding seems to be about these state laws that are above the federal law. Part of that, there could be -- certainly privacy advocates have certainly said that HIPAA's privacy rule is very weak and maybe that's why we're not seeing bad effects. So, maybe there's a trade-off between having a law that really works. But I think that there are other elements of the privacy rule in HIPAA that might be useful for states to think about. When they're thinking about setting their rules that are stricter, maybe looking at those dimensions in particular and maybe scaling back in those ways.

MR. CHAVEZ: So, I agree with the point about smart and careful regulation. But I will say that there actually is a significant amount of competition in the area of privacy. So, Microsoft, Yahoo!, Google, many other companies have been competing in the area of data retention, for example. You know, they talked about logs retention. Microsoft has launched Internet Explorer 8 which has some very interesting privacy features.

By the way, is this --

MS. ATHEY: They're recording, actually, so they want --

MR. CHAVEZ: Oh, are they? Sorry, recorders. Hey, where was my privacy notice?

**(Laughter.)**

MR. CHAVEZ: In private browsing in IE8. Chrome, likewise, has an incognito mode. So, there is actually a tremendous amount of activity at this point, I mean, frankly, you know, in a lot of ways. I hear this a little bit, I think, from the FTC and I see this in the FTC's privacy principles is that, you know, proceed with caution in the area of legislation because there really is a lot of thought, for example, being given to this notion of a value proposition for consumers when, again, the thought of maybe something like a discount on a product, if you're going to be placed into a particular category, you know, sports lover.

So, there is just a lot of stuff going on. I guess I would just say that it's worth monitoring and looking at and making sure that companies are really kind of living up to the promise that I think that we've presented here. But we should be careful.

MR. ACQUISTI: Well, the way I would slightly reframe the question is whether privacy can be a competitive advantage for firms. So, whether selling privacy can become a -- rather than just getting data, can become a source of differentiation, as product differentiation and practice differentiation.

The evidence I brought up earlier would suggest that, no, because people -- in the trade-off between privacy and money, they go for money. But we have other results that show, in fact, that under certain conditions, certain conditions, people will pay for privacy. The conditions are you need to show very clearly what the consequences of better regulation would be. The privacy alternative should be very easily accessible. So, you reduce transaction costs, cognitive costs and so forth.

In that case, people react to privacy as a form of feature that drives their choice of company, which means that, to me, because these conditions don't always happen frequently in many markets, a co-regulative approach is the best, one in which there is a basic background of legislation which protects some rights which cannot be transacted away. After that, the market can, indeed, allow for people to give away their data and so forth.

MS. ATHEY: So, just one last comment on the paternalism point. I guess I would just pause before saying that we need to paternalistically protect people’s privacy because, after all, for your average person who’s not a celebrity and not a politician, in fact, what are the objective risks to them from having this data shared. They’re fairly small and, in fact, you know, for most people, getting a discount on their Internet service or a coupon on the product is objectively worth more than whatever financial risk anyways they would be subjected to from the privacy. Again, the prevalence of medical identity theft is very, very small. So, you know, if we were going to be -- I would be very cautious to regulators about trying to paternalistically protect people against something that’s not really a significant risk, even as we need to make sure that some basic, basic principles are upheld.

MR. ACQUISTI: So, maybe I was unclear because I wasn’t necessarily advocating paternalism, but at most something which people call self-paternalism, which is considering cognitive biases and not making a decision for the user, but putting the user in the condition to make a better decision because the point is that no matter what you do, companies are making the decisions for the user. For instance, on Facebook, by choosing the default (inaudible) already decides what the user will do because we know from HCI research, human computer interaction, as well as economic research on 401K and many other fields, that people tend to like the status quo. Therefore, companies have a huge responsibility when they decide what the default settings of a system are.

We cannot just deny that the choices they're making are going to have a (inaudible) implication for -- sorry, an impact on how people will make decisions, therefore, on the end welfare of users.

MS. ATHEY: All right. So, should we take a question?

AUDIENCE MEMBER: So, we've heard some people say that there are some privacy laws that have created some inefficiencies. What are the sorts of privacy laws that some states have enacted that caused the problems and are there sort of obvious things that, as a policy person, I should look at that and say, oh, that's going to cause a problem?

MS. MILLER: So, I wish that I could point to a specific thing in the regulations, say remove that and we’ll eliminate this effect that we’re finding because there are certain common elements. The laws that we look at are laws that apply to hospitals and that have requirements for obtaining express consent whenever the information is shared at all.

That could really pose a problem for something like participating in a regional information network where the hospital that originates the data, the hospital that collects the data from the patient would not even be able to share with that network because they wouldn’t have control over future transactions. So, putting it in a common pool would basically be ruled out by a rule like that, that requires express consent for every time that data’s going to be accessed.

Now, there are --

AUDIENCE MEMBER: (Off microphone) (Inaudible).

MS. MILLER: So, the Federal Privacy Rule does allow blanket consent and, basically, HIPAA privacy rule ends up meaning for patients that whenever you go to the pharmacist for the first time or the doctor for the first time after the rule came into effect is you sign a form and that’s all it means to you. And if you want, you can read a long document explaining that you signed the form and that your rights are protected.

There are also rights there in terms of the access rights for patients to be able to see records and to challenge mistakes and correct the record. But this express consent, this authorized consent, every time the information is shared seems like a potential target. And then some of the rules don't necessarily say explicitly every time you use it, you need a consent, but they're kind of open or vague on that.

I think one of the dangers and one of the complaints that certainly has come up in the industry about these privacy rules is that there's a lot of uncertainty about what exactly is covered. So, there's sort of a fear of liability where hospitals may not even -- it may not even endanger them. You know, they may not be coming up against actual privacy laws, but some of the laws get worked out through cases and people don't want to be the case that gets settled for multi-million dollars to figure out that you broke the law.

And the other complaint is about a patchwork. So, some of these hospitals are parts of systems that span multiple states. So, there can be a cost associated with trying to accommodate or trying to comply with different laws in different states. So, that's another area of sort of low-hanging fruit where maybe there's benefits to coordination.

MS. ATHEY: Other questions? Scott?

MR. STERN: So, maybe this actually builds a bit on Amalia’s -- what are the regulations. But the general question I had was, you know, if I thought about the phenomena of privacy and, you know, people talked about the Facebook generation, my sense if that a lot of people -- and certainly anyone under the age of 30 -- sort of believes that Bessie has left the building, right? That everything is out there, that there’s a tremendous amount of personal information out there already and that closing the barn door at this point is too late, so that the marginal returns they face from doing anything active is extremely, extremely low.

In some sense -- so, I guess my question is -- I think the reason people believe that, and it was alluded to, I think, in the HIPAA regulations where everything is quite different than IT, is it would be very difficult, I guess for me -- I hadn’t -- when I’ve thought about it, to figure out what people know about me or the devices that I use. In other words, how do I know what information Google has about me, not in principle with the general policy statement, or Microsoft or, for that example, the IRS or whoever it is, or Visa and Mastercard. How do I know what information people actually have?

I would think that that -- because people have so much uncertainty about the level, they're completely inconsistent in their preferences about the margin. I guess that just seems to me the kind of -- you know, you could imagine that policy would be very usefully constructed that would, for example, give people kind of an audit, you know, that would say, not just any one company in particular. This is not a private company thing, it's a policy issue. But as best as we can tell, here's what people know about you.

MR. CHAVEZ: It's fascinating. It's like one of the core questions. I'll just point out kind of a difficulty and potentially something that I might characterize as an irony of this.

So, in order for you -- so, for example, if you were to sign up for a Google account, there's this option, it's called web history, where you can actually keep track of the searches that you've done and you can pause if, for example, you happen to be purchasing your wife's birthday present and you don't want her to see, or you can delete. And the challenge, though -- so, that is a feature that basically gives you access rights and correction rights. But you actually end up knowing a little bit more about you. So, you actually have to log in. You actually have to identify yourself to us.

AUDIENCE MEMBER: (Off microphone) (Inaudible) even if I do everything that Google (inaudible) I still feel like, boy (inaudible).

MR. CHAVEZ: Like the big picture?

AUDIENCE MEMBER: (Off microphone) My sense is that people have a very (inaudible) notion. In fact, I think (inaudible). (Inaudible) people have (inaudible) because they believe that the information out there about that (inaudible) whatever. It's actually much, much bigger than it actually (inaudible). In fact, as we saw (inaudible) very limited information for a limited amount of time (inaudible) with respect to the individual. But, in fact, people believe (inaudible) do you think that, you know (inaudible) knows everything you're doing? Basically, 85 percent of the people would (inaudible) yes and they're checking it all the time.

**(Laughter.)**

MS. MILLER: But they might be worried about the government knowing everything they're doing, too.

AUDIENCE MEMBER: (Off microphone) I guess what I'm saying is (inaudible) somehow (inaudible). In other words, do you (inaudible) that you essentially have accessibility of how information about you can be used (inaudible).

MR. ACQUISTI: If I may (inaudible) on this, I think it's an absolutely crucial point. It's just not --

it’s not just an issue of incomplete information. It’s also an issue that even if you protect yourself, everything, you shred all your documents, encrypt your email, anonymize your browsing, check all the possible information about yourself any time, still you’re not really protecting yourself. There could be, say, a company like ChoicePoint that you’ve never heard about who knows your SSN, your date of birth, your mother’s maiden name and sells that information to criminals. That happened around three or four years ago, the famous ChoicePoint case.

So, my point meaning that there is almost rational (inaudible) that people realize that we cannot know all the information about ourselves out there. We certainly cannot protect ourselves from the misuse of that information, so what’s the point in trying to stop this avalanche of data?

There is another related issue which is a privacy externality in that the more people who reveal more about themselves, the more costly it becomes for you to remain private. Not being on Facebook nowadays for a teenager is like not having a mobile phone for us or not having an email for a professional. Your cutting yourself out is a significant cost of privacy.

MS. GLUECK: I think it might be interesting to take a look at -- in the EU. Are people -- where people have a right of access under the national privacy laws to contact companies and say, hey, what data do you have about me, and if it was provided, say, as part of a registration experience, the company is legally obliged to turn over that data. Are people actually doing that? Are they requesting corrections or is it merely limited to, you know, that occasionally you update your email address when you get a new job because you want to continue to receive the newsletter that you were enjoying while at your old job?

It would be interesting to see, you know, has that really proven to be an important right for individuals or is no one exercising it at all?

AUDIENCE MEMBER: I was just going to comment from another area. In the area of credit reporting where today your credit report and the score that's implicit in that credit report affects not only your ability to get loans, but it also affects the price you pay for auto insurance and the price you pay for homeowner's insurance and whether an S&P 500 company will hire you and all kinds of other things.

So, as the information got to be more important, there were burdens put on those companies to allow consumers to get credit reports so that they could verify that this information that was now quite valuable and collected in a very central place could be verified.

The thought of every Internet company that collects data based on cookies having some kind of reporting obligation, I'm going to make you write that rule.

**(Laughter.)**

AUDIENCE MEMBER: And control the costs.

MS. GLUECK: You couldn't authenticate the users in that just based on a cookie ID. You'd only know that it was that -- probably that particular machine. But --

AUDIENCE MEMBER: Assuming it hadn't been hijacked by somebody else.

MS. GLUECK: Exactly. So, there's no way -- you couldn't actually provide access to the data because there's no way to authenticate it.

AUDIENCE MEMBER: But I think the layering that has to go on in this discussion is what information is important enough to trigger a burden and all information is not alike. I mean, the fact that you're searching for a car is not particularly troublesome, your health information might be more sensitive, what goes into your credit score and is that accurate financial information because it affects so many decisions that you make and the prices you pay, that's different quality of information. So, I think part of what has to go on in this next generation of discussion is when does information get to be important enough that there needs to be some policy response.

MS. GLUECK: I think that there's an impact. I mean, how much of what you do online would be embarrassing offline.

AUDIENCE MEMBER: Yeah, well, I keep giving my sons that lesson.

**(Laughter.)**

MS. GLUECK: For many years, porn was the only thing on the Internet that was making any money. So, people do -- they ask questions about, they do search queries about things that it might be embarrassing for them to have revealed to their neighbors.

MS. ATHEY: Having looked through search logs, I will say that, yeah, the eight-word porn queries are generally embarrassing.

**(Laughter.)**

MS. ATHEY: And there's a lot of them.

AUDIENCE MEMBER: (Off microphone) How long (inaudible) to actually (inaudible) rely on those (inaudible) agencies that are also part of (inaudible)?

AUDIENCE MEMBER: Yeah. I mean, it’s a much simpler problem because there were three companies that were the core credit reporting agencies and everybody consolidated around those companies.

MS. ATHEY: All right, we should probably wrap up, so thanks very much, everyone.

**(Applause.)**

**(Day 1 concluded.)**

**C E R T I F I C A T I O N   O F   R E P O R T E R**

MATTER NUMBER: P085800

CASE TITLE: ANNUAL BE INDUSTRIAL ORGANIZATION CONF

DATE: NOVEMBER 6, 2008

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the notes taken by me at the hearing on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: NOVEMBER 25, 2008

______________________________

ROBIN BOGGESS

**C E R T I F I C A T I O N   O F   P R O O F R E A D E R**

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

______________________________

ELIZABETH M. FARRELL