FEDERAL TRADE COMMISSION

I N D E X

FEDERAL TRADE COMMISSION

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

Thursday, November 4, 2011

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

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

P R O C E E D I N G S

- - - - -

DR. McALVANAH: My name is Patrick McAlvanah. I'm a staff economist here at the FTC. I'm legally obliged to go over the morning announcements again for the newcomers today.

The conference is going to be recorded by a stenographer, and so please try and use the microphones when you're speaking. And when there's questions afterwards, we'll be handing out microphones for the questioners, so please try and speak into those.

We are going to be collecting evaluation forms. They're the little yellow forms floating around. So, just before you leave, feel free to turn them in at the front desk.

The restrooms, if you leave here and then go slightly to the left, there's a black sign, and you can follow that.

There is Internet, wireless access, in here. There's a pamphlet at the front table that will tell you how to log in.

And then the security briefing is if you leave the building after you've already checked in, you do have to go through security again. It's just a minor hassle, but just be aware that it will take a few extra

minutes if you are planning when to come back for a session.

In the event of a fire, just evacuate the building out the front, and we can usually congregate in front of the Georgetown Law School across the street. In the event of something where it's safer to stay inside the building, then just basically follow the crowd, but we'll evacuate down into the parking lot. And if you spot any suspicious activity, just please alert the security desk right up at the front.

Okay, now for the fun stuff.

So, it's my pleasure to introduce Nancy Rose. She's our next keynote speaker. Nancy is a professor of economics at MIT and the director of the NBR's research program, Industrial Organization. She received her BA in economics from Harvard and her Ph.D. in economics from MIT, and her research focuses on the empirical analysis of firm behavior and the economics of regulation.

Nancy?

KEYNOTE ADDRESS

DR. ROSE: So, I want to thank -- start by thanking Joe Farrell and Chris Adams for inviting me to be part of this scientific advisory panel. It's been a really fun process, as you can tell from the papers that you've seen. The quality of submissions in response to the call for papers was really tremendous. I want to thank Patrick for helping me with the session that I was responsible for organizing yesterday afternoon. And just to thank the FTC staff in general. I think this is a tremendous conference. I've missed the last couple because of conflicts that my husband had, having to be out of town the same time. His presidency at the National Tax Association seemed to trump my just "I'd like to be in the audience at the FTC Micro Conference," but this year, I got to trump him. So, he had to wait until this morning to fly down for his NTA. He's someplace in Washington now.

Anyway, it's a great -- it's a great event, and as I was thinking about what to talk about, I was in the process of preparing a paper for a -- an AEA session that's a memorial for Fred Kahn, who passed away last winter, just before the AEA meetings, who some of you may recognize the name, has been sometimes affectionately called "the father of deregulation." So,

I've been thinking about Fred, gone back to reread some of his work.

I'm also editing a volume on regulation for the National Bureau of Economic Research, and I'm, as all of you are, continually bombarded with discussions about regulation or, more specifically perhaps, the evils of deregulation and how that has caused us to come to this terrible position that our country and the global economy is in at the moment.

And so at the risk of preaching to the choir, I've decided that I would like to talk a little bit about regulation today, and in particular, about putting IO back into regulation. So, I will preface this by saying I'm probably not so much talking to those of you who are on the FTC staff or the staff of other regulatory agencies. I think you guys have got it. So, I'm talking more to those of you who still sit in universities and academic research centers, where I think perhaps there might be an argument that we've moved a little too far in thinking about regulation as economic history.

So, first, just to set the stage, as I said, I was taking stock of regulatory economics. It's been four decades since publication of a number of classic works. The two that I would particularly note are Fred

Kahn's two-volume Economics of Regulation, Principles, and Institutions, which came out in 1970, and George Stigler's Theory of Economic Regulation, which is enjoying its 40th anniversary this year.

This year also happens to be the 25th anniversary of the Laffont and Tirole JPE article that really, while it's not the first article to address this issue and certainly not the last, is maybe the classic, if we had to pull one out of that extensive literature, that forces modern regulatory design to confront the problem of asymmetric information.

And then on the policy side, which is, of course, particularly important and relevant sitting here in Washington, it's been over three years since we really kicked off an amazing wave of policy reform, which I would date back to the -- its beginnings to the 1978 Airline Deregulation Act. And in that wave of policy reform, we dismantled a lot of the regulatory apparatus, much of which dated back to the Great Depression.

And I think what I'm concerned about is that we are perhaps close or at least discussing reinstating a lot of regulatory apparatus in response to the Great Recession, that while it may create employment for economists 20, 30, 40 years down the road, I think

wouldn't be a good place for the country to be.

So, let me first just stipulate -- and if you're interested in the -- I haven't given you an extensive bibliography documenting this, but if you're skeptical and interested in this, let me know, and I would be happy to share with you maybe even the reading list for my graduate economics course on regulation.

Those regulatory reforms, kicked off with the 1978 Airline Deregulation Act, have yielded really substantial, extraordinary benefits. We've seen increased productive efficiency and lower costs across a number of sectors. So, everything from improving the operation of generating power plants to eliminating empty back hauls in circuitous routing and trucking, to increasing airline productivity, and while we may not all like the increased load factors, it's certainly made the airlines more productive. And, in fact, if you think about current concern with energy use, a lot of these reforms in the transportation sector have had really dramatic effects in terms of improving the energy efficiency of that sector.

So, lower costs, lower or improved efficiency, lower prices in many sectors, which sometimes, although not always, was the point of deregulating or restructuring. You know, regulatory reform was meant to

improve allocative efficiency and operational efficiency, and that didn't always mean lower prices, but in a lot of sectors, it has meant that. Improved investment decisions and risk allocation. I'm going to come back to that in a little bit. And really quite extraordinary innovation gains.

So, one of the things that we know, looking backwards, has been a challenge is that if we require government agencies to be very involved in approving what firms -- how firms operate or what products they offer, that it's really hard for regulators to be able to be as innovative and nimble as firms might be able to be.

And so when we think about static costs of regulation -- and I'm thinking particularly here about economic regulation, but we could move to other forms of regulation as well, I think this is going to be a challenge as we think about financial services -- understanding that tension between innovation that increases welfare, social welfare, and innovation that's maybe solely or primarily for the purpose of evading or escaping some of the regulatory constraints is going to be a real challenge.

Now, the -- you know, regulation was not -- regulatory reforms were not all rosy and wonderful, and

in particular, they involved a lot of redistribution among stakeholders. So, you know, consumers' gains might be producers' losses or union employee losses. There were not just gains across stakeholder groups, but redistribution within stakeholder groups, and some of that has created lingering unhappiness with regulation, with regulatory reform, that I think is part of why we see it popping up repeatedly, is people arguing it's time to go back.

So, where are we now? And this is what concerns me, and this is why I want to make this pitch to my particularly academic colleagues. I see a current kind of popular regulatory credo that attributes deregulation -- and it's usually called deregulation even in sectors where it's really quite laughable to think that we've deregulated, right? So, if you look at financial services, there is probably no sector that is subject to more regulation by more agencies at a variety of different levels than the financial services sector. It's true we have changed regulation in that sector over time, but it is hardly deregulated, but we call it deregulated or the press calls it deregulation, and even some distinguished economists call it deregulation. But that's a current cause of our current woes, that markets are imperfect, and imperfect markets can be improved

through government intervention, that we know how to regulate better, even if we haven't in the past, and we just need -- and here, the list depends a little bit on who's talking -- but clearer legislation, more specific instructions, and smarter, harder-working, more honest, more public-spirited regulators, which I would say to those of you who are regulators, I don't think that's generally been the problem. That will not come as a surprise to you, although it might come as a surprise to those reporting these kinds of stories.

Unfortunately, this -- you know, we're all used to reading reports in areas that we're -- that we work in in the media that don't seem to quite exactly nail what we think the issues are. What concerns me most about this is that many of these articles of faith about regulation and regulatory reform are shared by a number of often distinguished economists. And so, the discussion yesterday about financial literacy had me thinking. The problem might be, in part, that we don't have regulatory literacy, that I would venture to say that there are very few industrial organization courses in U.S. universities that spend very much time at all, if any, on understanding regulation. That I think there are probably perhaps none, almost none, in business schools, and I'm not even as worried about the MBA

courses, because I would be a little concerned that as competitive strategy courses tend to try to help the MBAs think how do you compete to increase your profits, and maybe it's not always about competition, I might worry that regulation courses targeted to MBAs are helping them to be more creative in thinking about how to innovate around regulation. But I worry a lot about it with respect to particularly finance departments, because I think we are training generations of finance economists who are being asked to weigh in on issues of how to structure financial services regulation, who know very little about what IO economists know about the lessons of past regulation.

And so, I come to a place where I think this is maybe a really great example of Churchill's concern, that those who fail to learn from history are doomed to repeat it, and so here's my pitch. Here's some of the history, right? So, the first law of regulation Allan Meltzer wrote in 2009 was that lawyers and bureaucrats write regulations and markets learn to circumvent the costly ones.

And we could come up with a laundry list of examples of this, but this ought to be at the front of policy-makers' minds. This ought to be a reason why economists within agencies should be empowered, because

we tend to understand incentives. And this ought to be a lesson that we're teaching new generations of economists who are coming through.

Or if you don't circumvent, right, if you can't beat them, join them. So the other place that we see firms responding to regulation is by using regulations to create rents, all right?

So, one of my favorite examples -- I debated here, too. This is one of my favorites, so I decided to go with this, is free conference, right? Freeconference.com, fabulous organization, free conference calls -- not quite, because the people who are on the conference call have to call in a long-distance number. How do they make money? Well, there's collusion between the Iowa regulators and the Iowa rural telephone companies and Freeconference to kick up the access charge that they charge the long distance companies for calls terminating in the little Iowa communities and to split those rents between Freeconference and the rural teleco, right?

So, the example here was in 2007, when AT&T noticed that its termination charges went from $2,000 a month to $2 million a month in a community of 57 households, right? If you create an opportunity, you're asking for firms to figure out how to make money off of

it. It's taken the FTC until this fall to be able to finalize an order addressing this kind of compensation, and there is continuing litigation. One of the court cases involving Quest is scheduled for appeal hearings in early December.

My other example is one that's maybe near and dear to the FTC, which is the Hatch Waxman Act, which was intended to increase the ease that we could get -- to increase generic entry into the pharmaceutical market when patents expired and created this wonderful road map that allowed for collusion -- I'll call it collusion -- between -- cooperation -- between a branded pharmaceutical company and a generic entrant, where a generic firm could sue, claiming that the patent was invalid, and the branded manufacturer and generic plaintiff could reach a settlement agreement that effectively blocked all generic entry into that pharmaceutical through -- and the generic manufacturer agreed to that, because they received a set of payments as part of the licensing agreement, right?

So, I know the FTC was involved in raising this issue, Jeremy Bulow, when he was here, wrote on this, and it's been a challenge where we have a very admirable goal, but firms, if we create opportunities for rent creation, find ways to use that in ways that perhaps we

didn't intend.

So, how can IO economists help? I'm going to run through this pretty quickly, but a set of all Is, of what economists need to know and what I think we need to convey in our courses and our research about regulatory economics, and my pitch is really for putting more of this back into what we teach -- what we teach students and how we're researching.

And I guess I should say, which I didn't earlier, I think economists -- and particularly IO economists -- should give themselves a hearty pat on the back for the pivotal role that they played in the policy reforms of the late seventies and early eighties and nineties, because policy-makers pulled economic studies off the shelf to show, convincingly, that they understood what the costs of current regulatory policy were, what the effects of the reform would be, and then to document that, indeed, when we instituted those reforms, you know, most of our -- most of the work that we had done before to suggest what those effects would be was accurate. There were some surprises, right, but that economists created the work that folks in Washington needed to do their job.

And what concerns me now is that apart from energy economics and a little bit in health, that the

economists have largely abandoned the field of regulation. I should say, I -- or IO economists have largely abandoned the field of regulation, and so here's some reminders of things that we need to go back to thinking about.

So, information. Who knows what? I think the theory of asymmetric -- of regulation under asymmetric information has been a huge win for the world. It has transformed regulatory policy over the past quarter century. Just the simple acknowledgment that firms generally are better informed than regulators are and are going to act on their better information and that regulators who don't understand that fail to understand that at their own peril.

The centrality of incentives and regulation, I think, owes a lot to this theoretical literature and to subsequent empirical work that's been done. But this has really been a theory contribution and revolution.

And I would say Ofgem in the UK -- I don't see where Mark is, but Ofgem in the UK is the poster agency for the various -- front row -- is the poster agency for this. They have been really incredibly innovative, very focused on this, on the role of incentives, and on adapting regulation, which has been really crucial, because you can put a regulation in place. There are

some surprises in the way firms respond to it, and if you just sit on your hands, you've got an ever-increasing problem. So, they've been quite innovative.

Unfortunately, there has been much slower defusion of the implications and regulatory design based on this in the U.S., particularly as we push it down to the state level, where some state regulators understand this, but there's still a lot of regulations happening in state public utility commissions, and many of them don't.

Institutions. So, I -- you know, I come from MIT, where one of the first things we learn is that economists can't just sit and theorize or do empirical work from a theoretical standpoint. You have got to understand the nitty-gritty of the markets that you're studying. And I would go back here to Paul Joskow's '74 Journal of Law and Economics paper as a fabulous explication of this. So, Paul plopped himself down or found himself in the middle of a literature on what was called the Average Johnson Effect, the distortionary effect of rate of return regulation on capital choices by, particularly, electric utilities, a huge emerging literature, lots of theory, some tests.

The empirical stuff was pretty murky, and Paul

stepped back and said, "Wait a minute, regulators don't set rate of return which continuously binds. Yes, that's an input to determining prices, but regulators set price, not rate of return." And that simple insight turns a lot of the implications of the AJ model on its head.

So, you know, understanding really what the institutions are, if you're going to do work in this area, or teach students how to model and how to do empirical work is crucial.

I would also just shout out -- call a shout-out for that paper because Paul identifies the centrality of consumer aversion to nominal price increases as a really important feature of that market, and you might call that behavioral economics before we had coined that label.

Industry structure, of course, for IO economists, that's our bread and butter, but industry structure interacting with regulation, I think, is something that's often underappreciated. So, sometimes we leave regulation to public economics, and in public economics, as near and dear as that field is to my heart, having a husband who works in that area, you know, they tend to think of either monopoly or perfectly competitive markets.

And as we all know, oligopoly markets don't look like just something in between those two. And so, understanding the role of industry structure and how that affects regulation and its reform and its design can be really crucial.

The graph that you see -- that you can barely see -- is a graph of prices in three different electricity markets. The big spiky one is California; that won't surprise any of you who know anything about electricity. While it is true that the California electricity market was more highly concentrated in terms of generation in the restructured market than Pennsylvania, New Jersey, the PJM market or the New England market, we might be -- so, we might be forgiven for thinking, "Oh, that's a problem, we didn't make that market structurally competitive enough when we restructured it." The Bushnell, Mansur, and Saravia paper in the 2008 AER points out that IO economists can't just think about horizontal market structure. You need to think about vertical market structure.

So, Mara and Silke have made a career of identifying the importance of vertical organization and vertical relations. You can do the same in electricity markets, and what that paper demonstrates is that while you might think that a -- it looks like a Cournot model

explains California pretty well, it's not behaving as a perfectly competitive market, and at first blush, it might look as though PJM and -- oh, I have two -- no, that's PJM -- PJM and New England look more competitive. That's just because if you don't think about what the vertical organization of the structure is, you think that -- that Cournot competition in those markets would imply much higher prices.

But it turns out that when the regulators or when the market was reformed in PJM, that the distribution companies held on to their generating units. And so, if the buyers of electricity are also sellers of electricity or producers of electricity and, therefore, much smaller net buyers, they don't have the same incentive to raise prices as a freestanding, generating-only company would.

And so, what the AER paper demonstrates is that vertical structure plays a very important role in making prices look more competitive. If you don't have the right vertical relationships, you're going to see much higher prices in those markets. And so, when we think about regulatory design in electricity markets, we want to think about both horizontal and vertical market structure.

Incentives, as I mentioned, I think are really

at the core of what IO can contribute to getting people to think creatively about economic regulation, and we've just -- there's, again -- right, we could have hundreds of examples of how firms respond to incentives. One of my favorite papers in this area is Greg Crawford's 2000 RAND paper on the 1992 Cable Act.

So, that Act imposed price caps, price caps per channel, on cable companies, and what Crawford shows is that if those price caps -- if you just took the price caps, there would be very significant welfare gains from reducing prices, but if you didn't -- but you had to also think about what cable companies did in response to that, because they didn't just set prices.

They also set service offerings. And since the price caps were implemented on a per-channel basis. If we think for 30 seconds, you might say, "Well, how would we expect companies to respond?" Well, if you increase the number of channels, of offerings, in your basic tier of cable, which is what's regulated, you get to charge a higher price.

Now, that might be a good thing for consumers. You're offering them more channels, unless you're putting in channels that nobody really wants to watch, right? And maybe, in fact, what you want to do is take out the channels people want to watch and put them in an

upper tier that's not regulated.

And so what Crawford's work demonstrates is that if you just looked at the mandated price increases, you would have welfare gains on the order of one to two dollars a month per household, but if you look at what happens, accounting for the quality change, the change in the redesign of packages, that the estimated actual impact is basically zero or, depending on how you think about controls like year fixed effects, maybe even a loss.

All right, interest groups. I don't have to say much about this in Washington, but to those of us outside of Washington, a reminder that we need to think hard about that. So, this is a map from a paper that Steve Holland and Jonathan Hughes, Chris Knittel and -- I'm not remembering Parker's first name -- have recently done. It's looking at fuel policies, and they're comparing the -- it's a simulation of gains and losers from a variety of different policies.

So, the top one is cap and trade, which doesn't exist. The other three all exist, at least someplace. A low-carbon fuel standard, which says the average carbon content of fuel has to meet some standard. California's been thinking about that. A renewable fuel standard, think ethanol, one of our favorites -- I'm

sorry, that's a renewable fuel standard, and then direct subsidies for ethanol and biofuel.

And the thing to take away from this picture is that the more dark kind of masses you see, those are really big gains from that policy. So, what you see in this picture is that the low-carbon fuel standard, the renewable fuel standard, and the subsidies programs all have these pockets, these concentrated, high winners from those policies, all right? The light colors, the very lightest are losers, because these policies are all costly in terms of economic effects. They're not thinking -- they're not measuring the externalities. So, right now, they're just taking the direct winners and -- or the direct costs and benefits -- and profit benefits or consumer benefits.

And under those three that have got those masses of dark, most of the gains are going to producers, particularly for renewable fuel standards and subsidies, all right? So, we see kind of concentrated benefits and pretty diffuse losses, and that's a recipe -- so, Patrick mentioned I had an economics degree from Harvard. I actually had an economics and government degree.

I started out as a political scientist until I decided economics was giving me the models I needed to

deal with a changing world. But my courses in political science taught me that if you wanted a recipe for where you might see government intervention, it was when you had a few winners who gained a lot from a policy, and the losses were spread so smoothly that the losers -- it wasn't worth it for them to organize an opposition.

And so, it might not surprise us to look at this and say, "Gee, maybe this helps us to understand why we don't see cap and trade, the economist's favorite solution to this, in opposition to the low-carbon fuel standards or renewable fuel standards or subsidies, even though the costs of those alternative programs are so much larger, particularly the costs per unit benefit."

So, the left graph is the marginal abatement costs per unit of carbon dioxide that you're saving from a cap and trade kind of policy, where you've set the CO2 on fuel -- this is looking only at fuel -- and implement it through that, and that left -- the one on your right is -- is a low-carbon fuel standard, right? So, much, much, much more costly. But we don't see cap and trade with any political traction, right, because the winners from low-carbon fuel or from ethanol or from renewable fuel standards are out there lobbying, and there's nobody arguing for cap and trade.

So, if we want to get cap and trade, it's going

to take a lot more work, I think, by economists -- or programs like that, it's going to take a lot more work by economists to show the enormous costs that we're incurring by moving away from those.

Imperfections aren't only in markets. Imperfections are also in any type of policy intervention. I guess the -- you know, again, no great surprise to this group. I'm sure I won't have quibbles here, but solving last year's problem is pretty easy. I think we could all solve the financial -- we could all avoid the financial crisis, looking back now, by smart design of regulation. But the problem is putting policies into place to solve last year's problem or two years ago's problem is not going to prevent, necessarily, the next problem, which we haven't even imagined yet.

Smart regulators and regulation need resources, and I think that's an area that we've really starved over the last two decades. We can't expect our regulatory agencies to do first-rate work with kind of third-rate resources. And it's really important that we consider trade-offs with eyes wide open. Even regulators with the best intentions may play what I call Whack-A-Mole, right? You push down fees here and they pop up someplace else. How do we think about that in

our design?

And then the final thing I'll leave you with is this idea. Innovation, right, I know this is the political buzzword of the day, right, but innovation can cover a multitude of sins. So, we may have imperfect markets, but if we can't -- if those imperfect markets are innovating faster than our slightly less imperfect regulated markets, we might prefer the market imperfections wildly, right?

So, we could quibble with Jerry Hausman's numbers on the precise welfare loss due to cell phone delay, but even if you quarter his estimate of 50 billion a year due to the delay in cell phones, we're still talking a lot of money from FCC regulation that made it more difficult to bring that technology to the market and delayed it for a number of years. And I think that's just an important thing to remember as we are thinking about where to come out in the policy debate.

So, I leave you with this thought. Is the past prologue? This was really terrifying to me to see an opinion piece in The Baltimore Sun arguing that it was time for Maryland to stop playing victim, and we had the solution, which was for the Public Service Commission to build new generating plants and tell the local public

utilities they had to buy power from those and we could reregulate Maryland electricity one generating plant at a time. No thought of why we got to where we got to and how this precise policy led to the calls for regulatory reform ten years ago.

So, my argument, for those of you who are training students or thinking about research projects, is let's not regulate economic -- regulatory economics to economic history. Let's reinvigorate it, let's make sure that we're sending students out trained in the tools of understanding how to think about regulation, how to design regulation, so that we're helping to supply the research and the students who can help those of you in Washington who are doing the job on the grounds, with their boots on the ground on this one, to do the most effective job that you can.

Thank you.

(Applause.)

DR. McALVANAH: Thanks a lot, Nancy.

Next we'll have a paper session, our final paper session today, it's going to be chaired by Mark Armstrong.

PAPER SESSION FOUR:

ECONOMICS OF ADVERTISING MARKETS

DR. ARMSTRONG: Hello there. I'm chairing this session, which is on three papers to do with the economics of advertising. I'm hoping, I'm expecting, all three papers to be very interesting.

Just to recap on the rules, it's 20 minutes for the author to speak and then sort of seven minutes for the discussant and three minutes for the audience. So, we've got to stick to that half-hour slot per talk. So, please bear that in mind.

So, the first one is -- yes, yes, they're all here -- Joshua Gans, and Glen Weyl would be the discussants. Thanks.

DR. GANS: Thanks, Mark.

This paper is co-authored with Susan Athey and Emilio Calvano. Susan is here. Emilio is in Italy enjoying the economy there.

The paper here is -- we originally had a title, "Will the Internet Destroy the News Media?" And for a good year, we presented a paper with that title. So, we've got the softened version now, to look at the impact of the Internet on advertising markets for news media.

Our motivation is the news, the news basically

that newspapers, in particular, are in financial trouble, and they've been hit hard, and the culprit is, nominally, the Internet.

The reason they've been hit hard is sort of indicated by this chart of doom. This shows newspaper ad revenue from 1960 to the present day, and you'll notice that around 2003, 2004, there was a massive collapse in that ad revenue.

The top line shows total ad revenue, including classifieds, and the line we're interested in is the middle one, which is the advertising revenue from normal display ads, not from classifieds. So, classifieds are a different story, but the stuff that -- the ads that were actually interspersed with the news itself has suffered a similar fate, and arguably, that is what was actually subsidizing the creation of that news. And you'll notice the bottom little sliver there, that's the online advertising revenue going through traditional newspapers for their online sites.

Now, the story here is fairly obvious. The story is that people have shifted away from the print consumption of the news to the -- to online consumption, and their advertising revenue hasn't followed them, and it hasn't followed in a dramatic way. It's something like 50 times less per reader than for online versus

offline. And so this is the call. This is the problem.

What have been the reactions to this? Well, first is to try and nod it. It is basically saying, "Oh, dear, what explains the low ad rates for online is that online advertising is just ineffective. You move an ad from paper to a screen, and the consumers -- it doesn't work at all." That's our theory. It hasn't been borne out by the studies that we know of, especially coming out of the marketing literature. My colleague, Avi Goldfarb, is responsible for some of them.

I would say, if anything, unless you do something ridiculously stupid, like flash up ads in front of people and cause them not to be able to read anything, advertising is just as effective, if not more effective, online. So, that's not going to explain that catastrophe.

The second reaction has been, "Oh, damn, advertising is no longer there for us, the traditional news media model, where we would compete for customers and then rely on the ad revenue for our source of profits. It is broken. We need to do something else. We, namely, need to find other revenue sources." And so, it's basically, this is broken, the industry is stuffed. Unless we find other sources, who knows what

we're going to do? And a lot of the focus and attention has been on finding those other sources.

So, basically, this is a version of what is called the water bed effect. I think, Mark, you came up with this. I don't know. I've seen it in your --

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

DR. ARMSTRONG: But basically, if one side of the market's revenue disappears, then you've got to make it up from the other side, just as you would the contours of a waterbed.

Now, this is related to things of interest right here, various policy reactions, okay? The first is the obvious one. "Well, we've lost all this stuff. News is good. News was funded by advertising. That's gone. News is still good. Therefore, news organizations need to find other revenue sources, so they should be able to erect pay walls." But it's kind of hard to erect pay walls. If you put up your price, you see, your customers go away, and that doesn't help. So, what they want is for all the newspapers to put up their prices together. So, they come to the FTC, maybe we can all erect a pay wall together that will work.

Or we'll go to the people of the news aggregators, who are sort of referring people to sites, but they're also taking maybe some of the ad revenues,

and we should be able to go to the news aggregators and negotiate licensing fees. Now, if one newspaper goes to them, they'll say, "Get stuffed," but if all of them go together, as a bloc, then they'll be fine. So, basically, they would like a license to collude. They need to collude because competing isn't working.

The alternative is, "Well, if you don't want us to collude, how about you just hand us money?" And then there's a final thing regarding and a whole discussion surrounding tracking of consumers, which is another part to this. I have a little bit to say about that here, not that it's obviously necessarily a bad thing, but something which we don't quite understand the costs and benefits, okay?

So, our hypothesis is, you know, if you look at that chart, what the Internet has done is it's done something to the advertising markets, and we should try to understand that before thinking about all these other policies that are due to fixing other parts of the newspaper business. So, we're going to focus on that, and that's basically what we want to do.

And one of the reasons we want to do that is as soon as you think about advertising as a market, you think about what? Supply and demand, okay? And as you think about supply and demand, you think, what is the

problem here, okay?

First of all, in order to supply ads, you need people to view ads. You need them to pay attention to them in some dimension. That attention is still scarce. It's still scarce. So, there's not been some massive expansion in supply, in the ability to serve up ads to people, because people still don't have as much time to look at the ads. So, that's still something that's going to keep ad prices high, you would think.

The second part is the advertisers are still extremely interested in accessing that attention. So, the demand side of the market hasn't changed at all. Now, some people may argue at this point, well, the consumers are actually -- their entire attention has gone elsewhere. They don't consume the news anymore. But, in fact, when it comes to news media that is read, people actually reading the news, that has actually been growing for the last decade at 8 percent a year, above and beyond. There was a decline through the 1970s with television, but actually, the Internet has actually spurred actual read news. So, from the point of view of that part of journalism, the advertising market, both supply and demand, has some health associated with it, and we want to start there.

What I want to convince you of in the very short

time we have -- and that's what the paper is about -- that you only need two facts together to put together a theory of why we've had such a massive decline in advertising markets, and those two facts get you a theory. But also, the model that gets you there points you into a number of things that we need to look at regarding the operation of advertising markets to understand what's going to happen now, what's going to happen in the future, a guide for empirical analysis to disentangle the quantitative effects, and also, a framework for policy.

So, the two facts are this: The first fact is the Internet has been responsible for a lot of things. The one we want to focus on is that it allows people to consume a greater variety of news. People would previously get a newspaper, linger over it, spend time reading it. Now, you're going to Web pages, you see a whole variety of news from different outlets. That's fact number one.

Fact number two that you need to believe is that, in fact, it's very hard to track individuals if they move across outlets in terms of their behavior, okay? If someone reads a newspaper from cover to cover, you know which ads they've seen during that whole process, because you've laid them out for that, okay?

But if they're switching between sites, delving in here and there, so on and so forth, you might be able to track them as they stay on your outlet, but as they shift between them, that's an entire -- another matter, okay?

So, for switching, it's kind of obvious. I don't think I need to convince you very much. The browsers help people find more news content. There's lots of -- the fact that it's free makes it very easy to switch between outlets, and aggregators, social networks, and search have also brought a greater variety of outlets. This little graph here shows as you people use news in the aggregators more, the Herfindahl of their concentration across different -- of attention across different outlets decreases and decreases by quite a bit. So, they consume more news through this.

In other words, the Internet has facilitated switching. In other words, the consumers who might have been single-homers are more likely now to be multi-homers. What does that mean? Well, the traditional model of media economics, you have a platform, like The Boston Globe and that those two things, it provides content, which they might be able to raise some money from, and then it sells on the attention, part of the attention that they grab from

consumers, to advertisers for even more money, okay?

When competition comes in, which is one of the things the Internet has facilitated -- but we had it before -- that competitor will supply content and also serve up advertisers. Now, how the market balances itself out depends on what the consumers do. In the traditional model that newspaper provides, from Rupert Murdoch on, espouse, is that what consumers are doing is we're competing for their attention. They're either going to consume The Boston Globe or The Washington Post, okay, they are going to single-home. And then once they've made that choice, the newspaper outlet itself has a monopoly over access to them, and so it can charge advertisers a monopoly rate.

So, while Zipcar here may, you know, in a monopoly world just advertise in The Boston Globe, if The Washington Post comes in and grabs some of those consumers, it just divides up its budget between them, but the actual price change doesn't -- it pays doesn't change, because there's still monopoly access, okay?

So, now, that competition does have an effect. The advertisers multi-home, the consumers single-home, the newspapers want them to, to grab their attention, so they might drop the price to them. So, there might be a balanced effect, but the pot of gold is the advertising

revenue, and that's not changing with competition. And there's several models that bear this out.

Anderson and Coate look at a model where they're really focusing on the degree of advertising that's going on, how many ads are served up to consumers, and consumers in that model find those annoying. So, actually, as you get more competition between outlets, outlets want to annoy consumers less to attract them, so they actually contract ad supply. So, if anything, that would increase ad supply -- prices that we see.

Ambrus and Reisinger had a model where the consumers didn't -- some of the consumers actually multi-homed, and they're the ones who were kind of indifferent between the two outlets and wanted to pick and choose. And so they had a model there where, again, you were choosing your ad supply to be less annoying to consumers, but now, actually, when these consumers were multi-homing, they were kind of beating down prices to advertisers. So, what you want to do is to drive them away by having more annoying ads, so that you only had your really captive customers. So, it could go the other way.

Now, our contribution is different. We're very interested in modeling very precise -- you know, in a more general way the consumer behavior, away from

single-homing and certainly not to pure multi-homing, but something stochastically in between, and we do that in the paper. Some consumers will be multi-homers, some will be single-homers. The point is the outlets can't quite easily tell.

But we're very interested in what that does to the change -- to the advertising revenue and also to the advertising behavior, the advertising behavior in trying to manage the situation. And our contribution is to take away the model whereby on the advertising side of the market, consumers come attached with revenue and think of that as just disappeared and, instead, embed a formal market in that process. So, that's the innovation of this paper to the true side of markets literature.

So, that market can look like this. This would be advertisers -- would be some sort of advertising demand and advertising supplied based on attention. We're going to focus on matching as a core thing that's going on here, getting advertisers matched with consumers in an efficient manner. This gets us very sensitive to technology, which is where tracking comes in, and we have to be very careful regarding the allocation of consumer attention.

And so -- actually, I'll skip over this. The

model can explain a lot of puzzles regarding competitive behavior. So, anyway, the advertising supply, we imagine consumers are the following: There are two periods for them, the morning and the afternoon. They consume some media in the morning, they consume some media in the afternoon, and then they go shopping. And in the morning and afternoon, they can be served up some advertising attention, and they have a process which could be a logit model or something embedded in it that allows them to think about how they select outlets and where they get an opportunity, between the morning and afternoon, to switch between them.

What that gives rise to here is an advertising inventory for an outlet that is made up of switching consumers, DS, and loyal consumers, DL, because consumers might just choose to stay with an outlet. The loyal consumers who stay with an outlet consume twice as many ads as the switching consumers.

On the demand side, advertisers don't care when they impress the consumer. They just like to impress them once, and they associate a certain value associated with impressing consumers. Some advertisers really want to grab all the consumers because they've got a high value; others, less so.

The core of the model is this: There's an

impression game that advertisers are forced to play as these consumers may or may not switch between outlets. A consumer comes to outlet one in the morning and gets served up an ad. If they stay in that outlet through the afternoon, Starbucks doesn't want to give them another ad, the outlet knows that, and so they serve up a different ad for the afternoon. The model works well. You get the right matching of consumers to advertisers, and advertisers just pay for what -- for impressing the consumer.

The problem comes when the consumer annoys everyone by switching between outlets. They consume an outlet in the morning, they see a Starbucks ad. In the afternoon, they could, of course, see a different ad, in which case nothing's changed. But if outlet number two doesn't know what outlet number one has done, this could easily happen, okay? And this is a dilemma for Starbucks. Starbucks, in this situation, has purchased ads on both outlets but wasted some of them. They could avoid this by single-homing, just purchasing on one outlet, but there will be some loyal customers there they'll miss. And so they now have to play this game. It's more complicated because of the behaviors of consumers.

And so what happens here is advertisers end up

either consuming too little or too much, and we see this in the data. So, we have wasted impressions, because they try to target three to five exposures, and they're either doing a huge amount for very little. You could solve this dilemma in lots of ways. To stop consumers from switching would be a good idea. Rupert Murdoch would like to do that by erecting pay walls. That's one of the arguments.

If they couldn't track anyone at all, it would all be a crap shoot, and this wouldn't matter. You could coordinate in time, just advertise in the morning, and that will solve it, but my morning is very different from Susan's morning, I know that. You could pay per click, but actually the power of ads are in display, and you could also implement perfect tracking, which is something we've imagined in the paper as a benchmark, but we're far off doing for some market design reasons. And also, let's face it, for regulatory reasons.

So, you get these missed and wasted impressions, and that changes the willingness to pay. So, it's all going to be the action on the demand side. If an advertiser single-homes on an outlet, they could multi-home on an outlet, or they could multi-home throwing even more ads onto one outlet. If they single-home, they get loyals, they get loyals and some

of the switches. If they multi-homes, they get all the loyals across all the outlets, but they can't get all the switches, because the switchers are switching about and they might not see the ad they want. But if they pay for more and more ads, they will eventually hit everybody, and that's basically the market.

So, what you get is you get a situation where you've got single -- somebody's willing to -- an advertiser is willing to pay for single-homing ads, the normal advertising rate. When you ask them to put another impression on an outlet and multi-home, the value of that additional impression goes down because of these wasted impressions. They could then put more impressions on, but, again, there's the diminution return to that.

What that's doing is relative to a situation of no switching, where we have got total supply and total demand; in no switching, all the advertisers would multi-home, and you get a simple market price. When there are switches occurring, what happens is the multi-homers' marginal impression price value goes down, and so some single-homers actually come to the market, they've got lower valuation, and there's nothing wrong with what they're doing. They're paying what they get for as a single-home, and there's no impression problem,

but because they've got lower valuation, you have effectively got a reduction in demand. And so, as you get an increased number of switches in this market, the price goes down. There's the comparative static.

There is some nuances that always come with this. If you get too many switches, the really high-value advertisers want to purchase even more impressions on the market. AT&T purchased 95 billion impressions in this country last year. That's like 500 for everybody here. That seems like a lot, but why do they want to do that? They want to attract everybody.

As more switchers come into the market, it's theoretically possible that the amount of high-value advertisers in an inframarginal effect will start to bid up the price again, and the demand curve could rise. This only happens, however, if the ad capacity is high, and we have some doubts in the paper that we expressed whether that would be the case. But ultimately, if there's low ad capacity, you get a -- profits decline with the number of switchers in the market, certainly relative to a benchmark of perfect tracking. Yet with higher ad capacity, of course, this could occur eventually, and you actually get high-value advertisers buying too many -- so many impressions that it actually bids back up the price again, okay?

This has implications for mergers, but you have got to be sensitive to what the merger's going to do. It depends whether it allows the outlets to merge their tracking technologies or not, okay? If there's no ability to do that, because, you know, The Boston Globe and The New York Times are very different, then it's going to depend on price discrimination, which is a whole avenue that this paper goes down to. And price discrimination and the ability to charge different prices for ad campaigns, depending on how many outlets you control, is an ability -- is a way that you can exercise some degree of market power.

We have a whole lot of other implications, such as what is the implication on public broadcasters being allowed to run ads? What are the implications of blogs coming in, sucking up ad attention, but not actually -- attention, but not actually selling people ads? That actually reduces ad supply, that sort of behavior, and actually causes less of the problematic switching.

Switching is a problem when people view ads on one outlet and switch to another and view ads, but if they are going to blogs where they don't see ads at all or Twitter or something, that's not a problem. So, you get some counter-intuitive ads, such as that actually this sort of behavior -- blogs and public broadcaster --

can actually cause ad prices to rise, okay?

Now, back to policies. The nub of this paper comes -- and I can only advertise it at this point -- when you get outlets with asymmetric quality. When you've got asymmetric quality, you can actually get a premium -- you are capturing more consumers, that's one, but you also, we show, get a premium in the ad market. That is, you also get a positional advantage.

You have a better product to sell to advertisers, because that gives you less reason to erect pay walls, because that can actually destroy that sort of advantage, but it also may -- the tracking may reduce competition to consumers in certain ways because of the way in which lack of tracking causes outlets to compete more intensively to get a premium in the ad market.

Finally, there's actually -- and this is really interesting. There's an increased incentive to disaggregate and focus on reach rather than total reader attention. The traditional news model is we provide a newspaper with in-depth coverage so that everybody can enjoy it with all of the whole day, okay? But actually, the rate of return is higher if you can capture a small segment of the market for a small -- the whole market for a small period of time. The CEO of I Can Has Cheezburger, which causes all these cat videos, right,

he wants to make the entire world happy for five minutes a day. Our model says that's a good idea if you want to make profits relative to what The New York Times or The Washington Post is trying to do.

Lots of generalizations we could try to do. I'm sure Glen will outline some of them, and I'll also caveat, this is just a theory, and there's more going on there, but it is a start, at least, for this literature.

Thank you.

(Applause.)

DR. WEYL: I want to thank Josh and Susan and Emilio for giving me the chance to read this interesting paper and to Mark for asking me to discuss it.

So, this paper, I think one of its main contributions is to raise an important puzzle, which is, you know, we've observed that the Internet seems to have greatly reduced ad revenue. And at first, you know, glance, it might seem that, well, that's just what, of course, it had to do. But they point out that it's not that obvious given the standard models we have, because consumers, while reading news, are still captive. And it should be possible still to sell that attention, and, you know, it's a bit surprising, given that the amount of time they spend reading news has not gone down dramatically, and we have gotten this 50X reduction in

the price that you are able to charge for that attention.

That doesn't seem to fit very well with the so-called traditional media economics models, such as those of Anderson and Coate. And this paper argues that a primary explanation for that may be the fact that switching by consumers between outlets reduces the value created because companies will accidentally hit consumers with the same ad multiple times, and that may make the marginal value of each impression lower than it would be if they knew how many times the consumer had already been hit with the ad. And this is based on the idea that beyond some point, the ad doesn't generate as much revenue per impression as an initial ad would or as an ad at the optimum number of times hitting the consumer would.

The second effect that the paper explores is the idea -- and this could actually go in either direction -- that the switching may not just reduce the value generated by those impressions, but also, for sort of IO competition reasons, may reduce the amount that can be extracted of that value by the news -- the news outlet. And as with most IO logits of this type, whether the switching increases or decreases the amount of value that can be extracted is extremely sensitive to

how you set things up. Depending on the nature of competition and the nature of price discrimination and so forth, you can get many different types of results.

So, they also show that there's a bunch of complexities relating to a number of issues that arise in these settings, such as how you track the viewers within and across outlets, asymmetries between different outlets, competition by nonadvertising or advertising public broadcasters or Web sites, pay walls, advertising timing, contracts, a bunch of different complications that Josh sort of ran through towards the end.

But it seems to me that sort of key question in the paper is whether any substantial fraction or a substantial fraction of this 50X decline could be accounted for by the redundancy of these impressions. It seems to me that, you know, Josh, at the beginning of his presentation, said that what we need to explain this 50X reduction is the two facts, the fact that people are switching and the fact that, you know, the marginal impressions, when someone is hit with the ad multiple times, may be less valuable. And so it basically, it seems to me, that the key equation is this one, you know, the revenue that you earn as a news outlet is basically the value created by your impressions times the percent of that value that you're able to extract.

And as I pointed out, we know that the second term is really hard to, you know, figure out very precisely how that's affected by any of this stuff, because we know that depending on how you write down the model of competition, you can get lots of different things going on there.

Now, that doesn't mean it's not interesting to get at that, but I think all of us would pretty much think you'd need to have a very rich structural model of how the competition works to have any chance at really getting a number there. Any applied theory exercise is just going to show that this is ambiguous and we can pretty much predict that ex ante.

So, it seems to me a sensible approach might be to focus on the first term and say, okay, you know, how much less is an impression worth if we've already hit them with it a few times. And, you know, Joshua points out there's marketing literature on whether the Internet or news media or print media are more effective in reaching people. It seems like you could do a similar exercise for the number of people who hit someone with an impression, and, in fact, I wasn't quite sure, but sometimes -- I have seen the presentation on this paper before. It seemed like there might already be data on that.

And even this, you know, you have got these graphs on how many times people are, in fact, being hit with the ads. So, we should, you know, be able to multiply, you know, the number of extra impressions times the reduced value of a marginal impression and figure out whether that's anything like 50 times reduction in the value. And, you know, I don't know whether the extraction is going to go up or down as a result of this. It seems unlikely that it goes by an order of magnitude.

And so, it seems like a pretty reasonable approach to figuring out whether this is a major causal factor, would be just to do a simple calculation like this. And my feeling is that this might -- in answering Josh or trying to substantiate Joshua's claim, be more useful than an applied theory exercise is going to be, because the basic idea that value creation goes down is just -- you know, you can just write that out as the multi -- you know, the product of the marginal value of an impression and the number of times people are being hit.

So, there's a bunch of other issues that are related that might be worth looking at, that I think are somewhat secondary, though. I really did have trouble understanding exactly what the relationship was to

Ambrus and Reisinger and the Anderson, et al., paper, because if the main point of the paper is just to make the point that when you hit people a second time with an ad, it's not as valuable, these were discussed in those papers. And so, I think it's important to emphasize what the marginal contribution is here.

I also thought that there had to be a bit more thinking about what else might account for a 50X, you know, reduction. And there was a little bit of discussion of how ads online are different from offline, but I think it could benefit from more of that.

Price discrimination in the IO side obviously plays a huge role, and whether, under tracking, there would be more or less price discrimination and what optimal price discrimination would look like, I think, is of a lot of interest. I also found it a little bit hard to compare the model to the Anderson and Coate benchmark, because there wasn't really a generalization of that, and there wasn't some case of the model which corresponded to that, because the quantity dimension was taken out. And so, in terms of comparability to the literature, I found it a little bit hard to figure out.

And I also think, you know, it's important, in general, to think about -- which goes a bit contrary to the standard rhetoric coming out of the news community,

not just, you know, are you guys making profits, which, sure, is fine for you but is perhaps not in society's interests versus, you know, what is your marginal incentive to create different types of content, which has what social value? And I think that would be an interesting thing to explore.

Anyways, thanks.

(Applause.)

DR. ARMSTRONG: We don't have time, but did you want to take a minute to respond at all or --

DR. GANS: (Off mic.) Oh, no, that's fine. And those are useful comments.

AUDIENCE SPEAKER: (Off mic.) Let me just say, I mean, just in terms of putting it in context of the literature, you know, the other two papers aren't really focusing so much on the switching consumers, and so if you think of this as the model of -- you know, we have got the partial multi-homing of the consumers leading to endogenous partial multi-homing on the part of the advertisers, and that's a really important factor when you look out at the strategic behavior of the firms in this industry. And so, these issues of reach versus depth and all can be very well understood there, as well as some of the issues in terms of what ad exchanges are trying to do, the alliances between firms, and so on.

So, it fits very well with the kinds of -- the model is capturing what the firms in the industry are struggling with in terms of first-order (inaudible) policy issues.

DR. ARMSTRONG: Thank you very much.

Let's move on. The next talk is Charlie Gibbons, from Berkeley, and Jidong, who was here earlier -- oh, yeah, there he is -- will discuss. Thank you very much.

MR. GIBBONS: All right, terrific. Thank you for having me here today, and I thank Jidong in advance for his comments as a discussant on this paper.

So, what I'd like to talk about today is ad server and firm strategies and contextual advertising auctions. So, I think the first important thing to do in this paper is explain what this title means. So, by ad server, what we mean is some organization, some firm, that puts together a list of ads to display on a Web site along some kind of online content.

And when we talk about firms in this context, we are going to be thinking about the firms that are advertising in that list, selling some kind of product to consumers. And one of the interesting things about the contextual advertising auctions that are a bit different from the advertising structure that Josh was just talking about is contextual advertising is really

looking for consumers to commit to a product as they're seeing the advertisement. So, instead of AT&T just looking to get your attention and doing that in the morning or the afternoon, you go shopping for a cell phone at night, contextual advertising is really meant to hit you while you're shopping. So, it changes your decision as you're out in the market looking for a product.

And these ads are put together by a generalized second price auction, typically, and that's how the ordering in this list is determined.

And so what we want to do today is we want to look at a few different things in this paper. So, the first thing that we have to do is we have to think about how consumers respond to these advertisements. And once we start thinking about how the consumers respond, we can think about, well, what's the firm going to do in response to the consumer behavior? What are their bids going to be based on how these consumers are reacting to the advertisements?

And one of the things we're going to sort of start off by positing is that consumers search these lists from the top to the bottom, and when we look at the number of clicks that sites get, this is typically the type of behavior we observe. But one of the

questions that the paper asks is, is this actually rational for consumers? Does it make sense to go from the top of the list to the bottom based on what the firms are doing in their equilibrium strategies? So, is this an equilibrium that can be sustained and make sense?

But the real part of the paper that I want to discuss today is given that this is rational on the consumer's part, given these equilibrium strategies by firms, what incentives do the ad servers face? In particular, we want to think about what changes would the ad server want to make in terms of matching the ad to the consumer and what he's actually looking for. So, changing the probability that an ad is actually relevant for a consumer in his search.

Now, we can think about how the incentives for an ad server to reduce search costs might come into play. More closely related to sort of how the ad server can change the product market, we can think about how the ad server determines the length of the ads listed are displayed. So, we might think about does the ad server want to display many ads or just a few?

And then, lastly, one thing that is going to be relevant, I think, in terms of the elephant in the room of an antitrust potential challenge to Google is

thinking about what incentives an ad server, who actually has a firm that's trying to advertise to consumers, has to boost that firm to the top of the list, sort of despite the fact that it might not be the best match for consumers' interests? So, Google raising up gmail to the top when someone searches for email, says, as opposed to maybe Yahoo or some of the other online advertising, you know, options -- sorry, email options.

And so, of course, the questions are, why do we care about this? The first is we just want to think about what sorts of incentives the ad server has to innovate, to improve their matching algorithm, to create better matches between consumers and firms. And then, as I already alluded to, we can think about how this relates to competition policy.

So, again, when we start thinking about a potential challenge to one of these online ad-serving firms, we've got to think about, you know, where's the harm? You know, essentially, they're giving these ads to consumers for free, and so, of course, free has to be in scare quotes for us as economists. We know that nothing is free. And that's going to be the crux of how there could be harm in this sort of market, and we can think about what that might look like.

So, just to kind of sketch the model so we understand the overall structure, we won't go into solving equilibrium patterns or anything like that, but we've got a unit mass of consumers that are going to be shown this advertising list. There are j firms in this market, only m of which will actually be displayed to the consumer. And the consumer actually has a kind of interesting set of preferences.

So, they have a kind of lexicographic-type preference going on. So, what they're going to do is a consumer might say he's looking for a sweater, and he'll say, "You know what, if I find a sweater that I like, I'll pay $50 for that sweater, but if I find a sweater that is, you know, a color that I don't like or isn't -- you know, it's a crew neck and I want a V-neck, I don't like the fit, something like that," they say, "That sweater, I'm not willing to pay anything for it."

And what we're going to do is we're going to say that the chance that a consumer likes the product and finds the product at firm j is relevant to him has a probability qj, and so we're going to call that the relevance of firm j to consumers in this market. And we won't have a market segmentation story or something like that, where some set of consumers are more likely to be matched by a particular firm than others. We're just

going to say that every consumer visiting a site has the same probability as far as both the consumer knows and the firm knows of qj, of finding a product they like.

And then these valuations have a distribution in the consumer group, and the reason why we have this set-up is to allow for product differentiation among the firms, to ensure that we don't have consumers that are just -- everyone's buying from the first firm and then no one's looking at the rest. So, this is a way to give a positive market share to every firm on the listing.

And so, one of the first big assumptions that we're going to do is we're going to think about all firms charging the same price, just as a baseline to start with. And the idea that you can think about here, the price is going to be set outside of the model, so you can think, if you want to tell a little story, that there's some competitive market for sweaters and the prices are set in that market broadly, and then, you know, these are the prices that consumers are going to face in this ad listing.

So, if we think about firm j, it has this relevance, this chance of making a match with consumers of qj, and it's going to have a margin of mj, and let j, this index that we have for the firm, be the rank of this full expected margin, I call it in the paper, of qj

times mj. And one of the things that we show -- I won't talk about it today -- is that this index is going to match exactly into the ranking of the firm in this advertising list. Of course, that only is going to hold for the m firms that are shown, but we can make m arbitrary. We can make it the size of the full market if we'd like.

So, the sort of behavior that's going to go on here is we're going to say a consumer comes to this Web site, they see this advertising list, and there is some chance they are actually going to enter the list and start clicking on the ads, and we will say that will just have some probability as zero. And, again, the consumer is going to start at the top of the list, and he's going to go to the first site, and he's going to have a look at that product, and he's going to say, "Is this a product that I like? Is this a product that's relevant for me?" And if he says yes, then his next question is going to be, "Okay, well, is my valuation above the price for this product?" And if he says yes again, then he makes a purchase, and his search is over.

But if for any reason he doesn't make a purchase, then he's going to go on to site two, again, with some probability. And so we'll call that S1, the probability of searching beyond site one, and this

procedure is going to just keep going down the list, and consumers are going to, of course, be dropping out of the market because they give up searching or because they find a product that they like.

So, one thing that isn't going on here is consumers aren't searching for the best price, which isn't really a problem in this model, because we're assuming that all firms have the same price. So, we're kind of abstracting away from searching for the best deal on a product in this model.

And so, with this model in hand, we can go out and we can actually calculate some of the relevant quantities for this market. And the first one is going to be the click-through rate, and so that's just the fraction of consumers that visit a particular site. And so, of course, that's a big -- a big topic for people that are advertising in these contextual auctions. Then, based on that click-through rate, we can determine the demand, how many products does firm j sell. What are their total sales?

And then from those sales, we can actually ask, "Well, how many sales do they make per click?" And the reason that's going to be interesting is because firms are going to be paying per click in these auctions, and so what's going to matter to them, when they're thinking

about how much to pay, is how much of a margin they expect to make on a click. And so that's the next quantity we can make -- we can calculate, is what's the expected margin on a click. And this is going to be the relevant quantity when firms are thinking about how much they are going to bid in these auctions.

And then, again, that naturally leads into how much firms are going to bid. We can calculate that in equilibrium, and from there, we can just calculate the total ad revenue created for the ad server. And, again, we're not going to go through those today. The focus is going to be in the last half to use these results for how much ad revenue the ad server is able to bring in and see what incentives they might have to kind of change the structure of this market in the ways that we talked about at the beginning in order to increase their revenue and how that squares with the desires of these firms and the consumers that are participating in these ads.

And so, we're not going to go through the math today. I think it will be more helpful to kind of look at some simulations and go through the intuition. And so, what we'll think about first is imagine that the ad server can increase their chance of making a match. So, basically, they can boost all of the relevances of the

firms on the list. So, by improving the match algorithm, maybe helping consumers create better search terms, things like that.

And so, just for concreteness, the simulations that we're going to look at here, we're going to say that there are ten firms in this market, nine of which will actually appear on the advertising list, and then we'll just assume -- again, this isn't really going to matter, just for concreteness -- that consumers aren't going to give up. So, they're just going to search through all of the listings until they find a product that they like or just get to the last site, the ninth site, and not make a purchase.

We'll start with all firms having the same relevance, the chance of making a match of 20 percent, and then the margins are going to be what varies, from 0.1 to 1.0. And so that's just what's going to be giving us this ranking of the firms from top to bottom.

And so, what do we think is going to happen here? Well, the first thing that we notice is, well, if we boost the chance that a firm actually makes a match, then that's good for the firm, because the consumer visits their site, they're more likely to make a sale. So, consumers are more valuable. Their bids are going to go up in response to that.

But the flip side is that if they start making a lot of matches at site one and a lot more matches at site two, then you're going to get a lot fewer people getting down to sites three, four, five, all the way down to nine. And so, these firms aren't going to like the fact that they're getting a smaller pool of customers visiting their sites. And so, that's actually going to influence the total amount of revenue that they bring into the firm, and it's going to influence the total amount of revenue that they pay to the ad server. And so, these are our two sort of contradictory effects that we need to see how they actually weigh out in a particular case.

And so, the first thing that we'll look at is how firm revenues change in response to this 20 percent increase in the relevance, going from 0.2 to 0.24. And so, the first thing that we notice is that for the firms in the different slots, the proportion change in their revenue is only positive for the first set of firms on the list. So, for only the first four do they make more revenue when they have higher search/match probabilities.

And, in fact, when we look at firm net revenues -- and so, when I talk about firm revenues, I'm actually meaning net of their costs here, and when I say

firm net revenues, what I mean is net of the bids that they had to pay to get these consumers. So, net of the bids, we only see that the top two firms actually make more money when the relevance goes up. And every other firm is worse off in terms of how much they make when these match probabilities went up, which I think is pretty interesting, that this sort of innovation by the ad server would actually have a detrimental effect on firm revenues for these bottom firms.

And so, we can go and -- and that's the firm side, and let's go and think about what happens on the ad server side, and what we see is that bids go up, just as we expected, but they actually go up by more than 20 percent. So, we said that the relevance goes up by 20 percent, but bids go up by more than that. And the reason why they go up by more than that is because it's, you know, 20 percent, you know, increase in the chance of making a sale, but it's really bad if you start falling down the list, because you start seeing a lot more -- a lot fewer consumers than you would have before this change. And so, they get really scared of falling down a slot, and they're willing to bid more to make sure they don't go down that list. And so that's why the ad server can actually extract more than they were before.

And, in fact, we see that borne out in the changes in the total revenue raised from each firm on the list. So, we see that in six cases, I believe, the lower click-through rate is offset by the higher bids. And in this case, ad servers' revenue doesn't go up by 20 percent. It actually goes up by a little more than that. And so we might ask how are these properties -- how do these hold with different increases in the relevance? And we see that it's pretty -- pretty consistent property. Ad revenue is going to go up as you increase relevance. So, this is now looking at total ad revenue across different percent changes or proportion changes in relevance, and we see that it's increasing. But this little picture, the ad elasticity, what I call the ad elasticity, is just the proportion change in add revenue, divided by the proportion change in the relevance. And we see that it's increasing, but at a decreasing rate.

And for firms what we see is that total revenues are going up on the list, so this is total firm revenues across different changes in the relevance. But, again, what we see is that net revenues actually peak and then start falling, again, because the ad server is able to extract more of the rents in these cases.

And so, firms actually don't actually want,

overall, the relevance that they have at 0.2. They would actually like it to fall a bit. And so, the interesting things that we see from these little examples is that firms have a clear preferred value in total for the relevance that they wanted, that would maximize their revenues net of costs and of the bids. But even in this case, we're going to see that the tap-ranked firms are going to gain, while all the losses are mostly coming from the lower-ranked firms. And so we need to think about how we feel about that.

So, if it's just that you have the best red V-neck sweater out there and you're getting all the clicks and you're doing better, maybe we see this is fine, but if we are more concerned about firms further down the list, we're concerned maybe about how this might influence the competition in the market, then this could be, you know, a worrisome issue.

And the relevance that the ad server is going to choose is going to depend on the cost of increasing relevance, but it's, you know, quite likely that it will be more than the firms want. So, we see a little bit of a disconnect in these incentives between the ad servers and the firms. But in any case, consumers are unambiguously better off with better search probabilities.

So, just quickly in terms of how the ad server is going to choose the number of ads, what they're going to do is they are going to choose the number of ads, m, so that this ratio holds, the ratio between the full expected margin of one firm to the firm just above it in the ad listing. So lower numbers are going to mean higher on the list, has to be big enough, and this means the difference in these margins has to be small enough.

And so, one other thing that we explore in the paper is that the ad server doesn't like dispersion in margins. They want to have a narrow range of margins, because they're able to extract more, because the opportunity for bid shading is less, because one of the features of these auctions that's been well established is that the firms will shade their bids, quite typically.

So, the last point that I want to talk a bit about is, again, this issue of imagine the ad server has -- you know, wants a piece of the action in the actual product market. So, they've got a division that wants to appear on this ad listing, and the question the ad server is facing is, "Do we just want to let this auction go forward, you know, let this listing be natural, maybe our firm ends up in the first spot, maybe it ends up in the fifth," or do they want to say, "You

know what, we're going to essentially subsidize our own firm and say let's make sure it lands at the top of the list." So, again, I think this is going to have important implications for antitrust issues, potentially.

And so, what are the benefits of doing this? Well, if you're at the top of the list, you get a whole lot of clicks. You get the first crack at the consumer market, more demand, more revenues, and that's good. The costs are going to be that you push higher-value firms lower on the list, and so you push down their bids. You make less money from them. And so you want to figure out whether the costs are more than the benefits or just the opposite, naturally.

And so, what we would expect is that the gains from this sort of behavior are going to be the lowest for internal firms that we would rank highly anyway. If you're going to be in the first slot, then this sort of activity doesn't matter. The second, it's not going to make a big difference. And then, it's also not going to be that profitable if your firm really isn't a good match at all, because you're just going to lose out on all of these firms that are better matches, and it's not going to be worth it in terms of the lost bids.

So, again, we will do a little simulation, and

what we see in this picture is down here. Imagine this is the rank that the firm's internal division would be just if they went through and bid honestly. And then we look at the change in revenue if, instead of this firm being where it would be if it was bidding sort of honestly, at arm's length, what's the change in revenue for the ad server, taking into account both the total ad revenue and the revenue from this internal division in the firm, how does that change if we move from this slot up to slot one?

And what we see is that, of course, there is no change in profits for the first firm, and then it's positive for every firm, peaking in the middle, just as we sort of intuited, just starts falling. And it's going to bump up for when you take a firm that otherwise would have been off the list and then you bring it on.

And so, just to summarize, what we see is that in the case of increasing relevance, we see that the interests of the search or the ad server might be different from those of firms, and consumers generally like increases in revenue -- in relevance. The ad server is going to want to have fewer firms listed than both the firms in total and the consumers would prefer, and that might be something that we would want to look at. And, also, the opportunity for the -- for this ad

server to boost its own firm to the top might also change the competitive structure of this market.

And so, when we're thinking about where challenges on antitrust grounds go, I think that there's some evidence in this model for these latter two points being important situations to consider. So, thank you very much.

(Applause.)

DR. ZHOU: So, the research topic of the paper is very interesting, and I also agree with also that it's important to incorporate consumer search behavior into the study of position auction and (inaudible) the design problems. It's because that will not only help us investigate the optimal design of position auction, but also help us understand the consequence of position auction, okay?

And (inaudible), I will first briefly summarize the paper, and then we will come down to some details. And eventually, I will discuss where to put this paper in the literature, okay?

So, this paper takes Varian, (inaudible) as a starting point, and in various papers, he just assumes that top positions are more variable than positions down the list. And he also tries to endogenize the declining variable of all the other ad positions by introducing a

search market setting, okay? So, it's not that difficult to understand if a consumer buys a product and they leave the market once she finds a suitable product and the variant above the price, then fewer consumers are remaining as the search process goes on, okay?

And after that, he also characterizes the equilibrium bidding, the heavier firms, and that part is essentially the same as in Varian's paper, so I will not comment on that part, because he also didn't even mention that part.

And the final and maybe the most important part, he also investigates the ad server's incentive to manipulate the (inaudible), such as product variances and the search costs. So, that part perhaps is the most interesting part of this paper.

So, let me first comment on the search foundation. So, actually, they also took kind of reduced form away to model consumer behavior. So, consumers just to follow some exogenous stopping rule, which is a little bit different from what I initially expected when I read the introduction and abstract, okay?

So, of course, advantage of using a reduced form, that is simple, but it's also a little bit ad hoc, and some assumptions need more justification. And the

main concern I have with this assumption is that giving uniform price, why should -- the consumer could find out his variation below the price, should still continue to stay in the market and to keep searching? So, if I knew my variation will be below the market price, I should leave the market once I find out in my variation, or even in the beginning, this, they know their variations, they should not enter the market at all. But in the paper, of course, you assume that consumer will find their variation only after they find out a suitable product, okay? So, this is a little bit hard to swallow.

Of course, they also argue that -- in the paper that if we consider price dispersion in the market, then maybe this stopping will make some sense, but the problem is once we consider price dispersion in the market, then those consumers who find a suitable product and vary the product above the price, mean they no longer stop searching, because they also want to look for low prices in the market, okay? So, that's about the stopping rule.

And once in -- they also claim in the paper that a distinct feature of the current model is that the value per click will decline along the list. The main reason is because when the search goes on, a high

proportion of consumers will become those low variation consumers, and they never buy product eventually, right? But this result clearly relies on the assumption that those low variation consumers will still stay in the market, okay?

So, in some sense, intuitively, we may think of that low variation consumers sometimes should just search less in some sense. In extreme cases, they -- they don't bother to enter the market at all. So, I think they also may want to think more about this result, also.

So, here, my suggestion is that they also (inaudible) more serious search foundation, like Athey and Ellison's paper did, or just to keep the current reduced model as it is and use it for some potential empirical work, with a belief that the consumer search behavior assumed in this paper could be supported by data, and the people may become more tolerant to it when they see data. Of course, my belief may be wrong, okay?

And then some -- there's assumption about uniform pricing. So, this assumption, uniform price in the market, can be justified in extended model with actual pricing decision of firms if firms share the same production costs. This is true, for example, in the (inaudible) framework, because in that case, the model

setting is quite similar, okay? So, actually, if you consider actual firm pricing decisions with the same product costs, each firm will just act as a monopoly. So, we have kind of the paradox in the market. So, each firm is just charging monopoly price, so the price will be the same, but that seems not a good description of the real market. And this assumption is also hard to be consistent, if firms have heterogenous production costs, okay?

So, here, I think the author could try to develop a model of position auction with consumer search and (inaudible) price competition. Actually, in the literature, we still lack a model in this way, with effective price competition. And the most interesting part is the ad server's (inaudible) to manipulate some market parameters, okay, but here the problem is that when we change the market primitive consumer search behavior and the market prices may also change. Of course, this cannot be captured in the current model with exogenous search rule, but this makes me wonder how robust those results are if we consider an alternative section with active consumer search, okay?

And, finally, we are to put this paper in the literature. So, we already have papers purely about position auctions, and then we also have papers about

position auctions with a search foundation, and we also have papers which investigate the price completion implication of nonrandom consumer search, which could be -- caused by the author to add positions. So, I think the author should more carefully think about his contribution of the paper compared to the existing works, especially those position auction papers with a search foundation.

Okay.

(Applause.)

DR. ARMSTRONG: Charlie, do you want to come back for a minute or are you happy, Charlie?

MR. GIBBONS: (Off mic.) No. I think those are all good comments. I appreciate it. So, those are definitely some issues that I've already started looking into, and we'll continue to think about it.

DR. ARMSTRONG: Very good.

Well, I am going to try and keep on time. We'll have a slight change to an empirical focus with Minjae talking about advertising in magazines.

AUDIENCE SPEAKER: We are going to load Ginger's slides first.

DR. ARMSTRONG: Okay, very good. And the discussant will be Ginger Jin. Okay, so I will just wait for the slides to come up.

AUDIENCE SPEAKER: Do you want to come around?

DR. SONG: Okay, thanks for having me here.

This paper is about how to estimate platform market in two-side market, and for the empirical implication, I will treat the magazines as platforms to try to attract both advertisers and readers a make a profit from both groups.

Okay. In two-sided markets, the two agents interact through the platforms, and they care about the presence of the other group on the other side, and platforms account for these cross-group externalities in making profits.

There are many examples of two-sided markets. Payment systems is one very good example, where the credit card is platforms, and try to attract both the merchants and consumers and try to think about the presence of each other. The video game systems are other good examples, where the video consoles try to attract both the game developers and game players. The advertising in newspapers, magazines, and Web sites, also very good example, and this is the -- advertising and magazines is my empirical application here.

So, my paper brings two important features of two-sided market into structural model. The first one is that agents on both sides care about the presence of

agents on the other side, and the second feature is that the platforms charge two prices, one for each group.

And I focus on the cases where the platforms charge a fixed membership fee. The -- another interesting -- an important pricing is the usage fee, where the agents pay per usage, but I will focus on only the fixed membership fee cases.

And I consider two versions of two-sided market. One is the two-sided single homing, where the agents on both sides join only one platform, and the second version is a competitive bottleneck model where one side single-homes, but the other side multi-homes.

Then, I think about how to estimate this platform market given the platform level prices, and then I show you how to estimate the costs and recover mark-up from these costs. And then at the end, I will show you a merger simulation, where I allow the publishers to merge and have a different market structure.

Okay. There are numerous theory papers on two-sided markets. The most cited ones are Rochet and Tirole, 2003 and 2006, and Armstrong, 2006. As of last week, the row she and Tirole, 2003, and Armstrong, 2006, are reaching, like, thousand citations. My paper is closely related to the Armstrong 2006 paper, in a way

that I focus on the fixed membership fee case, and I look at these two versions of two-sided market.

There are relatively few empirical papers, but the number is growing fast. Two papers worth mentioning is Rissman's 2004 paper on the Yellow Pages and Argentesi and Filistrucchi's 2007 paper on newspapers. The way that my paper is different from there is in Rissman's paper, the application is the Yellow Pages where there is only one price charged by platforms. So, the consumers do not pay to get Yellow Pages, but only the advertisers pay to post the advertising.

In Argentesi and Filistrucchi's empirical study, they assume the consumers do not care about advertising. So, by this assumption, with this assumption, they can sort of stay away from the cross-group externalities affecting each other.

Let me briefly introduce the model. The two-sided single-homing model is basically the same as the standard demand estimation model that we know in empirical IO literature. The important difference here is that the presence of the other side agents here, this SJA for side A and SJB for side B, are important platform characteristics. And in addition to price variable, this SJB and SJA are another endogenous variables that correlate with unobservables.

And once making a distribution assumption on Epsilon IJ of both sides, we have this -- the market share equations that look quite familiar to many of you here. Again, the SJB and SJA enter here as a -- as platform characteristics.

So, for each platform, we have a pair of demand equations to estimate. For multi-homing model, the single-homing side is basically same as the previous version, but the -- for the multi-homing, I follow Armstrong's model. So, the -- each agents -- the multi-homing agents make the membership decision independent of the other platforms, so they visit -- their member decision for the one platform is just to compare the benefit of joining the platform and the cost of joining the platform, which is the fixed membership fee here, and if the benefit is larger than the membership fee, they will join this platform. And they make this type of decision for all the platforms in the market, so they can join as many platforms as they want as long as the net benefit is positive.

And this benefit is a function of how many -- the agents, this platform attracts from the other side, and this platform's specific quality and the agent type, okay? So, if the multi-homing agent's willingness to pay is high, then they will -- so, in many platforms, if

their willingness to pay is lower, then for the same number of agents from the other side and for the same quality of platform, they will join fewer -- fewer platforms.

And given the distribution assumption on the willingness to pay, we can -- we can write down this -- the market share function for multi-homing side. So, basically what it says is the agents -- the multi-homing agents with the high -- the value of alpha I will join the many platforms, okay?

The -- because of this cross-group -- the externalities, the price elasticity is not the same as just the first -- the derivative of the market share function. Why? Because any price change in one side not only affects the market share of that side, but it also affects the market share location on the other side. But that change doesn't end there. It also affects the market shares of the original side, okay? And that changes, also, the subsequent impact on the other side.

So, this effect or any, like, small perturbation in price has this ongoing location effect because of the cross-group externality. And we call this the feedback loop. So, because of this, I treat the pair of market share functions as implicit functions and compute the

full price elasticity that traces over this feedback loop using the implicit function theorem.

The estimation is very -- very simple and standard. It's basically GMM estimation. So, the key thing here is to find the good instrument variable that's correlated with SJB and PJA, so the market share of the other side, and price, but are not correlated with a demand (inaudible), and I will talk about my instrument of variables in a few minutes.

For the competitive bottleneck model, instead of having these two equations for the multi-homing side, we have this market share function. We have this market share function for the multi-homing side, but we do observe almost everything except for the platform-specific quality. So, given the number of people joining from the other side and the number of people joining from multi-homing side and price, and given the distribution assumption on the RFI, we can convert this market share function to recover this WJT, the platform-specific quality, and then we reverse that on the non-price platform characteristics. But, again, this is also the system GMM with endogenous variables.

And recovering marginal costs is to search for the cost-absorbing agent on both sides that satisfy the first of the conditions, but in doing so, for every

(inaudible), we have to compute this cross -- the full price elasticity and cross-side price elasticity using the implicit function theorem. So, it's a competition. This is a little bit more burdensome, but we have computers that can do this.

For my empirical application, I look at the TV magazines in Germany. The reason that I look at the segment of magazine is to justify the single-homing assumption. So, here the consumers choose one TV magazine to buy, but the advertisers can advertise in as many magazines as they want. I have quarterly information on the copy prices, advertising prices, and advertising pages and content pages and circulation, et cetera, et cetera.

The data are collected by this German -- the public institution equivalent to U.S. Audit Bureau of Circulation. And the (inaudible) here is that I use -- so, I know the publisher, and I have the information of the publishers' magazines in other segments. So, I use the same publisher's -- the average price and average market share in other segments, like a business and politics magazine segment and car magazine segment, et cetera, et cetera, and I use this -- the average price and market share in other segments as instruments for the market share and price in the -- for the TV magazine

segment.

The average TV magazine's price -- the copy price is about one euro, but the one-page advertising cost, like, 30,000 euros, and so the magazines make a lot of -- much more -- higher revenue from selling the advertising than selling copies. The one euro for the 100-page magazine is a little bit too low in the context of one-sided market, but this is perfectly fine in two-sided market. So, let me show you three tables, and I'm done.

In this table, so in OLS estimates, what you see here is the copy price is negative, and with IV, it becomes more negative and significant, but the magnitude of the price coefficient itself is not big enough to justify the profit maximization in one-sided market.

The ad price -- the ad page is positive, and it becomes more positive with instrument of variables, which shows that the consumers or readers of the TV magazines actually like the presence of the advertising. And this is sort of consistent with the -- my other paper with William Kaiser at IJIO that shows that the readers in the magazine market do not necessarily dislike the advertising.

And I also have the magazine fixed effect and magazine time effect here. So, using these estimates, I

recover the market power. So, in the left panel, I have -- I assume the one-sided market structure, and on the right-hand side, I have a two-sided market structure. On the one-sided market structure, the median mark-up for the reader -- for the reader side is about 62 percent. It's very, very high mark-up. But in the two-sided market, the -- actually, the cost is much higher than what we estimate in one-sided set-up. Actually, it's higher than price, so what we cover here is that the magazine is actually making a loss by selling copies of magazines, but they make a lot of money from selling advertising pages, and if you compare this to -- so, the market slightly goes down by moving from -- the mark-up for the advertisers slightly go down by moving from one-sided to two-sided market, because here, the consumers like advertising, okay? Consumers like advertising, so for the -- given the same observed price, the model sort of estimates the -- gives us the lower mark-up for the -- the appreciation of advertising.

But the overall picture here is about two-thirds of magazines, they make a loss by selling the magazines, but the -- the copies of magazines, but they make their profit from selling advertising pages.

So, then the merger simulation, okay? The

observes market structure is oligopolistic market structure, so I simulated two extreme market structures. So, one extreme is a single ownership structure, so every publisher only has one magazine. That's a single, what single means. The monopoly means one publisher has all the magazines in the TV segment, okay? So, these are two extremes.

So, I have the -- what I have is oligopolistic, so I move towards the single platform, and I also move to the other extreme. So, if we assume the one-sided market structure, obviously the prices of the -- the copy prices always go up when the market becomes more concentrated. When it moves from the single ownership to the monopoly ownership, all the copy prices go up. But in two-sided market, that doesn't necessarily happen. About -- so, yeah, I only -- the -- put the selected magazines, about 70 percent of magazines lower their copy prices and they increase the advertising prices.

So, for magazines that lower copy prices, they always increase the advertising prices, and magazines that increase their copy prices, they always lower the advertising prices.

So, the mergers in this two-sided market are not necessarily -- in magazine segment, at least in Germany

TV magazine segment, the mergers are not necessarily harmful for readers. What about the advertisers? Advertisers, about 70 percent of them face the higher advertising prices because of merger, the higher concentration of the market, but that higher prices are also compensated by the larger reader bases, okay? So, they don't like the higher advertising prices, but they like the larger reader bases that they can get in a monopoly market. So, the overall welfare change is ambiguous. It depends on numbers. And in some of the markets that I simulated, I actually found the case where the total welfare is higher.

So, in this paper, I bring two -- the important features of the two-sided market into structural model and estimate the model and recover the mark-up and did a counterfactual analysis, and it shows that the platforms charge below marginal cost for one side, but they make a profit from the other side. And it's very important to account for both sides, in our estimation.

Am I on mic?

DR. ARMSTRONG: That's perfect.

(Applause.)

DR. JIN: I will first thank Mark for giving me an opportunity to discuss this interesting paper. As you can see, clearly the paper is -- basically it

combines BLP methodology with a two-sided market context. In this process, it's emphasized two unique features for two-sided markets. One is there is a positive externality between the two-sided markets, and rates incentive for the platform to be big, because it's going to offer valuable size of the market.

The other -- the other's feature is that the platform will compete on both sides. So, they could compete on both sides directly, as both -- to be single-homing, or even if one side advertisers would decide whether to advertise on each magazine separately, the platform may still -- ends up competing for advertisers because they are competing on the other side of the market, and the two sides are linked.

So, the findings are very sensible, just as we already know, magazines tend to set consumer price below the marginal cost in order to expand their customer base; however, they earn large mark-up on advertisers.

And then a merger into a monopoly could be welfare-enhancing, because merger itself is sort of increased value for both sides by maximizing the externality between the two sides, and this could even result in a lower price for consumers and for -- sometimes even lower prices for advertisers, and so this could be -- I guess implication is that merger in this

context could be much less anticompetitive than the merger in one-sided market, okay? And these findings are specific tied with the features in two-sided markets.

I think the whole paper is very clear intuition. The empirical implication is enormous work. I really want to praise Minjae in all the exhaustive efforts. If you read the paper, it has models on two kind of competitions. It even offers simulations before the empirical estimations, so that's really, really a lot of work, okay?

And I wanted to summarize a little bit on empirical estimation. It basically has three parts. One is on the consumer demand. This is kind of typical BLP story that you have some logit transformed market shares on the left-hand side. You have some kind of terms into the consumer utility on the right-hand side. And that right-hand side could include number of advertisers in the magazine, could include the price for the magazine, okay? And the typical endogeneity on the price will be solved by some instruments, and the endogeneity on the number of advertisers could be solved by instruments as well, okay?

The other side of the market is advertisers, as Minjae has described. The advertiser side is estimated

by imposing assumption of -- I think this F is on the -- like the distribution on the value of advertising, okay? And then that's per reader, and then this tied with the number -- given number of readers in the magazine and given the price on advertising, this would decide how many advertisers would decide to buy the advertising at that price, and by inverting the first equation, he will basically estimate a second equation that's kind of a function of advertising demand on the magazine attributes, okay?

So, implementation, it assumed the F function, the value of advertising function, to be not normal with some mean zero and variance, 1.4, which I think is a little arbitrary. I haven't seen too much justification why you would choose those two numbers. And as far as I understand, this conversion, assuming that platform knows the exact form of the F function, so it's sort of treating the price and number of readership to be exogenous. I think this is a little sort of in conflict with the other parts of the paper, okay?

And then the third part is assuming every platform engaging in the Bertrand-style profit maximization, taking into account both sides of the market.

So, I would like to offer some comments on

advertiser demand first, and this is already -- I already talk about the F function seems arbitrary, and because it's -- the shape of the distribution actually would describe how demand -- how advertising demand is sensitive to price and the readership, okay? And this seems to completely assume away -- it's sort of imposing -- I think it's a very strong assumption on how the market of advertising responds to price, and this is also not accounting for that the price may account for some, say, demand shock in advertising market, and that should be addressed, at least by instrument variables.

On the -- another issue that Minjae didn't talk about in the presentation but sort of talk a lot in the paper is, as we know, this is for two-sided markets. There may not be a unique solution of market shares given parameters, okay, and he argues that this would not affect estimation. I'm not completely convinced on that, okay? And also, this seems to directly affect elasticity calculation in the merger simulation. So, for example, elasticity would require to know how the price would affect the market share, but if there's multiple solutions to that market share, I would like to know, like, how you select equilibrium, for example, and when you compute the elasticity.

Okay, and the same thing for the merger

simulation. In that process, we need to know how this price affects market share and, therefore, profit, and then derive for the optimal price. So, again, that's -- we need to know exactly how we choose the equilibrium there before we know how the profit is determined, okay?

I also have some comments on model choices and IV strategy. The empirical estimation assumes competitive bottleneck, which I think is a more appropriate model than the two-sided single-homing model, but I would like to see more justification on that; like, for example, do you see the same advertisers do multi-homing across magazines, okay, and is there any exclusive dealing in the pricing strategy to violate that assumption, okay?

And on the IVs, it's basically assumed that the demand shocks are independent across different segments of magazines, and this could be violated, let's say, if different segments try to target the same readers, okay, or the same advertisers try to advertise in multiple segments, if the products are not specific to TV but more to, let's say, Starbucks Coffee or something, just to try to reach the readers. This is not -- this could be introducing some common shocks on the -- across segments.

And also, the publishers, if the publishers own

multiple segments and they introduce, for example, bundled pricing in multiple segments, this could introduce some correlation between the segments as well. So, I think the assumption you put in here is reasonable given the context, but I would like to see some justification on that, okay?

It's not clear to me, by reading the paper, whether you have accounted for a publisher may own multiple TV magazines or there has been actually market structure changes over time. So, I assume you have accounted for that in the -- in the estimation, okay?

So, finally, there are some comments. In one paragraph of the paper, it mentions that consumers and advertisers actually end up having different quality rankings, the magazine. I think this is derived from the magazine fixed effects. So, consumers may prefer one magazine to the other, which means the first magazine would have a larger market share; however, this first magazine does now charge -- was not -- was not more demanded by the advertisers.

So, this seems inconsistent, if you're thinking there is a positive feedback loop between the two sides, then their view should be largely consistent, okay? So, your explanation in the paper is saying that maybe the larger market share publishers are not fully exploring

their power in the readership; however, in the whole model, you are sort of assuming that they're doing this optimally. So, I think this is -- seems like to be some confliction -- conflict there, okay?

I can see mathematically that the merger may lead to some lower advertising price, but I would like to see more intuition on that, because it seems like they should explore their larger readership by charging a higher price, just intuitively, okay?

And, finally, as we have seen in the first paper of this session, that platform may differentiate, especially when they own a lot of magazines, in order to sort of offer better targeted advertising if you see more sorting between different kinds of consumers and different kinds of advertisers into different types of magazines. And I don't think this model has addressed that, but I can't think of, clearly, how this would affect your estimation in the counterfactual, if it is in the data but not addressed in the model, okay?

But overall, I really enjoyed reading the paper. Thanks so much for giving me the opportunity.

(Applause.)

DR. SONG: Did you want me to comment?

DR. ARMSTRONG: Yes, that's fine.

DR. SONG: For this no IV for the multi-homing

side --

DR. ARMSTRONG: Can you get to the microphone?

DR. SONG: -- the -- in the single-homing side, what we assume, we actually fixed the distribution natural -- the parameters of the distribution of x on ij, okay? So, we always do this. So, it's a -- it's a ordinary utility, so we have to fix the location of the distribution to estimate it.

For the multi-homing side, so we have to fix the distribution otherwise, you know, this whole distribution is moving around. That's why I fixed the mean and the variance of the f at that value, so -- and I didn't do that arbitrarily. I look at the -- the profit of the publisher and sort of gave the -- pick up the number that makes the publisher in the market make the non-negative profit throughout the existence. So, there's sort of this empirical issue where I have to really fix the distribution of the f.

And on the single-homing side, we always do this by fixing the distribution of x on ij. Yeah, that's it. Thanks for your comments.

DR. ARMSTRONG: (Off mic.) Very good. Well, I suggest -- well, I would suggest we thank the three speakers for their very nice presentations --

(Applause.)

DR. ARMSTRONG: -- on their related papers on overlapping readership and just (inaudible). Thank you very much.

AUDIENCE SPEAKER: (Off mic.) Is there any FTC announcements?

DR. McALVANAH: There's just going to be a short morning break and coffee outside. We have to be back by 11:30-ish.

(Whereupon, a morning recess was taken.)

DR. BECKER: It's my pleasure and my honor to introduce Mark Armstrong. He did his graduate work at Saint John's College, Oxford University, and he's been a thought leader on some of the most important issues in economics, both in his editorial roles at The Review of Economic Studies and The RAND Journal of Economics, and through influential articles about price discrimination, multi-product pricing, access pricing, and the issue that we just heard about, two-product pricing in platform markets.

He co-edited The Handbook of Industrial Organization, and he wrote "Regulatory Form Economic Analysis and the UK Experience." I'm excited to hear what he has to say, and so I'd like to bring up Mark Armstrong.

KEYNOTE ADDRESS

DR. ARMSTRONG: Well, thank you for staying the course here. It's a great pleasure to be invited to come and talk to you. So, yesterday, we were -- we heard -- we economists were encouraged to work on consumer protection issues. They said it had everything going for us that we could possibly want, and so I was pleased to hear that, because it's something I've been dabbling in myself in the last couple of years, and as well as all the good reasons that Janet gave. One other good reason is that there are very few people doing it. So, it's very relaxing, working in a relatively young crowd in the field, unlike merger policy, where there's just everyone -- everyone ever.

One downside is that there isn't very much money in it, so we were told the antitrust people had nice clothes and things like that. I've done my best, but I -- you know, I'm probably just the same as all the other tribe of economists.

So, I'm going to give a talk about an aspect of consumer protection, and I tend to think of consumer protection as sort of falling into three kinds of stories, three kinds of remedies. You're trying to mitigate information problems prepurchase, that's something we've heard about already; you're trying to

mitigate unpleasant surprises postpurchase; and the third kind of issue is you're trying to prevent sales techniques which involve duress or undue pressure or anything like that. So, those are the three kinds of consumer protection policies, at least the way that I organize it.

So, I'm going to present a contribution to the third of those, about trying to give an economic model of high-pressure selling in a very stylized sort of way. It's -- maybe I'm departing a bit from the standard style here. I'm not sort of giving a keynote-style paper, overviewing any particular thing. I'm just going to give a sort of fairly nontechnical overview of a particular paper, which is the one here, which is joint with Jidong, who you saw earlier on, and I've generally done most of this consumer protection work with him, okay? So, that's the -- that's the topic.

It's possible some of you have seen this, but I'm not -- yes, I think -- anyway, I apologize if one or two people have seen it, but it is, in fact, pretty different from the version that I've been trotting around elsewhere.

I've -- one of the issues of being a so-called keynote speaker is you don't have a discussant, which is a loss for me, because the discussants have been so good

in this conference and a loss for you, so -- but in compensation, please, just put your hand up and interrupt as I go along. That will -- you know, there is not going to be anyone else summarizing what I'm saying. So, please just ask questions as I go along. I'll keep my eyes out for your hands, okay?

So, what exactly am I going to talk about? Well -- here, let's move down. So, I would say there has been relatively little work done in economics about actual sales techniques, okay? So, it's a little bit surprising given that's sort of what markets are largely about, but there's not much about actual sales techniques. And the thing that I want to look at in this -- in this paper is the particular technique where you force a customer to decide to buy quickly, okay? In particular, before this potential customer knows what other options are available in the market, okay? So, this is a well-known technique. I'm not so -- I am not going to set -- try and say what's going on in the U.S., but in the -- the European policy is quite clear on this, since they have -- whatever it is, 30 black-listed forms of selling. Normally they're all pretty bad, things like, you know, no threats of violence or refusing to leave the home or something like that, but one of them is not giving customers -- giving customers

the impression that the offer is available for a very short time and not giving them time to evaluate other offers in the market. So, this is something that's addressed by policy. It is probably not a very easy policy to enforce, but we can talk about that maybe. Okay, so, that's the particular kind of high-pressure sales technique that I want to provide some economic analysis of.

Okay, and there are sort of three things that a seller might do to try and make people buy quickly, okay? Something that we're quite familiar with, when we sort of try and think about job offers or things like that, an exploding offer, okay? An exploding offer, I am going to model it as a situation where the customer is -- a seller has knocked on his door maybe, some kind of door-to-doorstep seller or something like that, knocked on his door and says, you know, "I'm only in the area now. If you want this vacuum cleaner, this (inaudible), whatever it might be, you have to buy now," okay? So, that would be an example of an exploding offer, as I use it.

Okay, and there are various other examples documented in -- there is a paper that goes along with this, that are documented in the paper. So, you have a photography studio -- this was before the days of

digital -- telling customers that they have to decide what pictures to buy that they shot that day since their negatives are going to be destroyed. That is an interesting example. So, some of you know more about law than I do, but the law journal system often involves exploding offers to authors, okay? So, you submit a paper, maybe submit it simultaneously to a number of journals, and they come back to you and say, "We're going to publish this, but you have to decide now," okay, before you find out from a better journal whether they are going to let you come in. Okay, so that would be an example of an exploding offer. And, in fact, because it didn't perform very well, the -- a number of journals have colluded, if you like, to commit not to make these exploding offers. There's a public document about that.

Okay, so that's one thing that's a rather extreme tactic that's saying you're in my shop or I'm in your home or whatever way it might be. You have to accept now whether to accept my deal, okay? And you can't come back even if you want to.

A sort of milder version of that would be what we would call a buy-now discount, and this is just a -- instead of banning return, you just raise the price if they buy later, okay? So, I come around and I offer you

a discount if you buy immediately, okay? So, again, we've got a -- sort of anecdotes, if you like, in the newspaper. A car dealer offers an extra $500 off if you buy now, so that he says he can make his monthly quota, okay? You have to give some story for why this thing is in place.

A landlord offers a $100 reduction if the tenant agrees right away, and there's a whole book about an anthropologist's experience in a direct sales force, and he gives lots of examples of this kind of behavior, and, in particular, one of them is the Kitchen firm he was embedded in would offer a long-term quote for the particular bit of work, but a 10 percent discount if you just agreed to it straight away, okay? They called it a first-call discount or saving time of another meeting or something. So, that's, again, trying to encourage customers to agree immediately before they have a chance to see what else is out there.

Okay, those two things that I've explained depend very much on announcing the policy to the customer, okay? You go to the home and say, you know, they only have any effect if you say what's going to happen, that they can't come back or it's going to be much more expensive if they come back.

A variant of this that has a very different

flavor would be the idea of a surprise price hike, okay? You go to -- maybe you go to, say, your firm is going to get antitrust advice from one of five consulting firms. You go to one. They seem to offer a rather expensive deal. You go to the other four, and they're all conflicted or unsuitable or whatever it might be, and you have to come back to this original person, and, you know, maybe they've got an incentive -- they know, presumably, what you've done in the meantime, and they maybe have an incentive to rack up the price to exploit their monopoly power.

So, that would be -- I put surprise in inverted commas, because if we are in a rational, equilibrium world, these consumers are going to anticipate that, but it's not announced and it's not committed to at the time of the initial meeting, okay? They all have the effect of discouraging an onward search because of the disadvantageous terms that happen when you come back.

Okay, we are going to talk about two scenarios. I'll probably only talk about the first one, which is a very simple -- really, a very simple story, hopefully not embarrassingly simple, but it's heading in that direction. We just have a single seller. Somebody knocks on the door and offers something, okay? And these customers are going to have some uncertain outside

option which they don't know at the time the seller knocks on their door, and that's going to be the key thing, okay? They can go and find it out, but they don't know what it is at that time.

And then you can use essentially the -- exactly the same arguments to talk about an oligopoly search model, which is basically just the same thing, except that the outside offer is determined endogenously by rival offers in the market.

Okay, what's the key difference between this and any other of the thousand search model-type things out there? Well, we're going to think of our seller -- to make this story work, the seller has to be able to tell apart, to distinguish people who come first and then people who come back to buy later, okay? And in these sort of stories that I've been telling, that seems quite plausible.

The doorstep seller obviously knows whether he's been there before or not; people selling you insurance, home improvements, et cetera, et cetera; car dealers, all of that kind of thing. But it doesn't apply to anything like going to buy groceries, you know, a supermarket doesn't keep track of whether I come in and go and come back again or anything like that. So, it only applies to these kind of -- well, largely

face-to-face-type interactions. There may be some issues to do with computers recording whether you come and go, as well, but think of it sort of like a face-to-face interaction.

And I am going to try and persuade you that firms will often then have an incentive to discriminate against the people who want to buy later through these various techniques, okay? So, my generous introduction said that I worked on price discrimination, and I suppose this could be thought of as one example of that.

Okay, what are the -- what are the reasons why sellers might do this? Well, there's a strategic reason, just to deter onward search, to make people more likely to buy immediately, and that is clearly a feasible thing for these sellers to do, and this is going to apply when you can commit to your sales policy, okay? It only works in that way.

The second kind of reason is the more orthodox price discrimination reason, which is that if you know they've seen your thing, didn't much like it, otherwise they would have bought it immediately, but they have gone away and found even less good stuff out there, and they come back again, what does that tell you about what you want to offer that person next?

Okay, so let me just trot you through the basic

model. As I say, it's something you can pretty well fit on a slide. So, we have got a single seller. This is a Bible salesman or something like that. And its strategy, which it knocks on your door, is to give you a price for the product and, where relevant, what its policy is if you decide to buy later. They might not allow it at all in the extreme case of an exploding offer.

Consumers have some uncertain willingness to pay for this item. We call that u, and that's just going to be idiosyncratic, and I'm just going to say that the fraction of with u, because in p, I am going to call that the demand curve, q of p, okay? And it's not supercrucial, but the firm doesn't get to see what u is in this interaction. Okay, so that's all unbelievably standard, but the twist is what happens if the seller doesn't buy or doesn't buy immediately this product, okay?

Her uncertain outside option, I am going to call that v, and she doesn't know v when they first encounter the seller, okay? V is bigger zero, you know, so you can't be forced to buy. It's bigger than zero, but there might be something better out there, for instance, an alternative seller with a better deal.

Okay, think of u and v as independent of the

(inaudible), not crucial, but I'll explain how that comes in. There may be a cost involved in discovering v, you know, going to another seller or something like that. That's the search cost. Okay, and that's all I need to say there.

So, let's see when -- why and when a firm might wish to do this high-pressure sales technique, okay? Say there were no search frictions at all, just to make my story as easy as possible, okay? So, you can costlessly go and investigate the rival outside option. So, the two policies that we can think about are allowing free recall, that means that you can freely come back and get it at the same price as I offer you in the first place. That's the regular kind of sales technique, the nonduressed sales technique, okay?

Given that, the consumer is always going to see what else is out there, because it might be better than offered by the firm. So, you're always going to do that, but you'll come back if you find out that u minus p, your net surplus is bigger than your outside option, v. So, what's your expected demand from this policy?

It's just -- remember, q is the probability that u is bigger than something. So, q of p plus v is the likelihood, the probability the customer will buy at price p if the outside option is v, and, therefore, if

you take expectations over that with respect to v, you get the expected demand under this easy-going sales technique, okay? That's where that E of Q is. So, that's your pay-off. That's your expected demand if there's a seller doing that.

Suppose you do the hard sell style. You make an exploding offer, and then what's -- what's the consumer going to do? Well, he's risk-neutral. He's only going to buy if the pay-off at the firm is bigger than the expected outside option, okay? He's not going to come back, so there's no -- there's no utility from that. So, that means that the expected demand is Q of p plus the average value of v. So, it's just taking the expectation inside the demand function. Okay, so that's the -- that is the -- in a nutshell the comparison between the two sales techniques and, you know, Jensen's inequality is clearly the relevant thing to look at here.

So, what's going to happen, it just depends on whether the demand curve is concave or convex. If the demand curve is convex, then you prefer to do free recall, to have the expectation outside the demand curve -- this is for any price -- and if you prefer to do the hard sell if the demand curve is concave. So, that is the sort of essence of the story. It makes it

clear that it's ambiguous. It's just not something that you always want to do or you always don't want to do. It depends on the fine details of the demand curve that you're facing, and you can't get away from that.

Some of you will be thinking, what about commitment here? Is it really credible that these things are going to happen? Well, this red proposition here is still going to hold as long as there's still some fraction of vulnerable consumers, if you like, who believe this incredible claim. Suppose that the seller can't commit to this thing. If you did come back, he will give it to you. He will give you the vacuum cleaner.

But all you need is some fraction of people who are credulous, who actually believe the sales technique, and this argument goes through, just applied to that fraction of elderly people or whatever it might be. I'm not allowed to say that here, but the vulnerable people, yeah. Okay, so you can see that it's -- a firm as an incentive to do this even when not every -- if it can't commit.

Okay, you can talk about the price effect of this high-pressure sales technique. What -- does someone set a higher or a lower price when it does that? And in general, it's a bit more fiddley. You don't have

such a nice condition, because it all depends on the expected elasticity of these two demand curves, okay? So, it's -- you know, it's ambiguous in all the cases that you are likely to look at. It is going to be more expensive to buy the product when it's sold using this high-pressure sales technique.

Okay, and in that case, there's going to be two problems in this market. There's going to be poor matching -- if you think about it, given the same price, this sales technique is bad for consumers, because they sometimes could be better off elsewhere or, indeed, sometimes they might go elsewhere and would be better off if they stayed, if they came back. So, it's bad matching between products and consumers, and there may be higher prices as well. So, that's a double-whammy. So, that's the story there.

So, this is a bit -- yes?

AUDIENCE SPEAKER: So, here you are assuming that consumers do not know their outside option, the value of the outside option?

DR. ARMSTRONG: That's right.

AUDIENCE SPEAKER: But you have some consumers who are highly informed of this outside value, and some are not, so there's a selection problem.

DR. ARMSTRONG: So, that is exactly the reason

why the search model is the natural extension to this. So, the search model is they've seen someone first, and then when you go to somewhere else, and so that is your outside option is the first offer. So, in a search model, you take account of that, of the fact that a fraction of the population do know what their alternative is, and it doesn't affect these kind of results, but it's -- that's why you want to do the extra effort of the oligopoly version of this.

Okay, so this would give you some kind of argument for why you might want to make an exploding offer to an employee or potential employee, but it does depend on the fine details of the -- of the thing, okay?

Buy-now discounts, this is the milder version of the same policy and probably more common than a literal exploding offer, okay? So, this would be slightly -- you go to some dodgey electronics store or something like that and you try to buy a camera and they say, "Well, I'm going off shift in an hour. I'll get my bonus if you buy it now. I'll give you 10 percent off." It's that kind of story.

And it turns out that it's much more -- it's -- I think of it as just universal pretty well, that you are going to want to do that if you can, okay? The previous result had this rather strict concavity versus

convexity condition, which, you know, lots of things are in the middle of that. Here, the only condition is that the demand curve is low concave, which is a pretty mild condition, okay? And in that case, you do want to offer a buy-now discount, okay?

What's the difference in the two cases? Well, the difference between the two cases, why is it so much more common here than it is with exploding offers? Well, because there's an extra revenue effect, okay? Suppose you rack up the price for coming back, that will boost the number of people who buy immediately, as before. It will reduce the people who come back later as before, but you also get more money from the people who come back, because you've increased -- you don't just shut down the market. You just extract more money from them. So, for that reason, you get an extra kick.

Okay, and when you work out the examples, it is often the case that this form of price discrimination, both the prices rise when you have this high-pressure sales technique coming in. So, that's a very -- in my experience, it's amazingly rare to have a monopoly model where price discrimination forces all prices to go up. It's because of the extra frictions in the market.

So, in my last three minutes, I am going to do the third thing, which is the surprise price hike, the

third version of this. This is our consultant firm being nasty and racking up the price when the person is over a bow, okay? So, when does it want to do this?

Okay, it's exactly the same model as before. Suppose that there's no announcement about what it's going to do. The customer naturally assumes that the offer remains on the table if she decides to come back. The question is, does the firm then, in that circumstance, have an incentive is to raise its price to those customers who do come back to buy later?

Okay, there are three cases. So, suppose there were no search frictions at all. Then, if you think about it the right way, the answer is clearly no, okay? If there are no search frictions, everyone's gone on to look for the outside option, and some of them come back if it's better. That's exactly the same scenario whether you increase the price or not. So, there is no incentive to raise the price in that case. It's like saying you've come to a shop and a customer has agreed to buy my TV at $500. Knowing that, do I want to raise the price? And the answer is no. It's just the same as that.

The middle case, a bit blurry, but the third case, at the bottom, suppose there is a small cost of coming back to the monopolist, okay? So, I have to make

a call to get the salesman to come back or I have to visit the shop, whatever it might be. There's some cost to coming back and not buying immediately. Call that r. Then the answer is you always want to rack up the price when they come back, okay? And the argument is transparent, and it's very similar to the Diamond Paradox. So, that returning cost is a bit like a positive search cost in the Diamond model, okay?

So, just to run through, suppose that p is this candidate uniform price that the seller makes. The consumer goes away and comes back. Whenever that condition is, u minus p minus this new cost of coming back, r, is bigger than the outside option, those are the people that will come back. And, therefore, the seller can raise the price by r and not drive any of those customers back to the outside option, okay? So, they can surprise all their returning customers with a discrete price rise, and that will -- is bound to be profitable, okay?

And, in fact, if you think about it, the same argument applies, that there can't be any equilibrium returning demand, okay? If there's any equilibrium price that the consumers expect when they come back, even if it's much higher, they can't -- they can't be in equilibrium for exactly the same reason. The firm

always wants to rack it up by another r, just like in the Diamond Paradox, okay? So, if there's no commitment in this model, it's not as if these discrimination -- these high-pressure selling goes away, which is -- you might -- whoops -- which you might think, but, in fact, it amplifies the incentive to discriminate against returning buyers, and? And, in fact, you are going to be forced to make an exploding offer in equilibrium in this model.

I think I better stop there. The same thing happens in search, the same kind of results. There's a few more details, but I think I've been told I've run out of time. So, I'll stop there. You get a flavor of the kind of thing that's going on.

(Applause.)

PANEL SESSION TWO:

PERSONALIZED MEDICINE

DR. GARMON: Thank you very much.

Now, for something completely different, we have a very distinguished group of panelists, very fortunate to have a very distinguished group of panelists to talk to us about personalized medicine. First, Dr. Adam Clark is a scientist and policy advisor with Medtran Health Strategies. His research focuses on molecular diagnostics, personalized medicine, and patient-centric care. Dr. Clark has worked with numerous patient advocacy and disease research organizations, including LiveStrong, where he's served as the director of science and health policy. He's also served as a technology development specialist at the National Cancer Institute, administering programs in cancer biomarker detection technologies. While at the NCI, he also performed policy assignments in the White House Office of Science and Technology Policy and the Office of the Secretary of Health and Human Services.

From the FDA, we're very fortunate to have the chief economist with the FDA, Clark Nardinelli. Before joining the FDA in 1995, he spent many years teaching undergraduate and graduate economics at various universities, including the University of Virginia,

Tulane, Clemson, and the University of Maryland, Baltimore County. Clark's recent research interests include work on best practices for FDA cost-benefit analysis, integrating uncertainty into the economic analysis of public health policies, and evaluating the economics of policies to deal with addiction.

And finally, last, but certainly not least, Mark Trusheim is a visiting scientist and executive in residence at the Sloan School of Management at MIT. He's been a special government employee for the FDA's Office of the Commissioner and is the founder and president of Co-Bio Consulting. Mark's research focuses on the economics of personalized medicine, particularly the integrated quantitative modeling of stratified medicine development and commercialization. Mark, along with --

MR. TRUSHEIM: They can read the rest of it.

DR. GARMON: Okay. I will leave it at that.

The only thing I wanted to mention as well before we get into this, both Adam, Mark, and Clark will give short presentations, but when we invited them, we gave them two questions. First, probably most important, what is personalized medicine? How does it differ from the traditional drugs, therapies, treatment protocols that we commonly think of as medicine? And

second, is the existing institutional and regulatory framework set up to promote the development of personalized medicine?

And with that, I'll introduce Adam Clark.

(Applause.)

DR. CLARK: Okay. Thank you very much, Chris, and thank you all for attending here, and we're going to be doing a Q&A then afterwards? Great. Then I'll try to go through these slides as quickly as I can, because I think just interaction, different perspectives, will be most valuable to everyone here.

Starting with this first question, what is personalized medicine? There are many different definitions, and I'm actually going to talk about probably a broader definition than most of my colleagues would. But overall, it's referred to as getting the right treatment to the right patient at the right dose at the right time. So, when we're looking at it in the context of research and development, primarily we're talking about diagnostics and targeted therapeutics. And this is when I -- you know, years ago, about a decade ago, as I was coming up through the National Cancer Institute, how we traditionally thought of it. Within the past few years, the patient community has started to gain an understanding for it but view it very

differently and much broader than those just in the science community view it.

This is the -- one of the hallmarks of personalized medicine. This is -- this slide shows various types of breast cancer, and years ago, we found this protein called her2/neu that was expressed in about a quarter of these breast cancers. They developed a test that could screen for this, and from these bottom ones, you can see that have high expression, they can -- they would give the drug Herceptin to treat these patients. So, you had a diagnostic saying a particular alteration was occurring, and we had a drug specifically to target it to those patients.

Now, as we've moved out then to the patient community -- there's my former boss, Lance -- patients view personalized medicine as being about them. There's a big overlap in personalizing care, and I wanted to show just on the left, this is Lance's treatment summary. For those that don't know his story, he was actually diagnosed with metastatic testicular cancer prior to winning any of the Tours de France. They removed his testicle, gave him surgery to the brain, and he started on a round of chemotherapy called BEP. All of these treatments come with associated risks. What's interesting, the BEP, Bleomycin, BEP is a lung toxin.

When Lance found this out, he said, "I cannot be on this treatment protocol. Cycling is everything that I do." And he had to go through and find some other doctor, finally up in Indianapolis, who said, "All right, there is a different treatment that's out right now, VIP. We can get you on that to spare your lungs." So, Lance's passion has been we need to find ways to get this information into the patient's hands so that they can make the choices based on the quality of life that they want.

As there now is this integration, then, between patient needs and the developing technologies, I think we're seeing some challenging decisions that we're going to face in a regulatory environment. So, an article came out in 2009, this was on CNN, that for general Y women with a cancer risk, it's just a boob, and it was referring to young women who carry what the BRCA I or BRCA II gene mutation that makes it very high risk for them to develop breast cancer in their life. So, we're seeing an increase in the removal of both breasts to prevent this from happening.

Now, here, we have a genetic test, but we do not have a disease out there yet, and so -- or that has manifested itself, yet the consumer market is saying, "We're going to make this choice and have surgery based

on a positive result to this test." So, this is the odd world that we're living in now, that we have many different technologies integrating.

So, PricewaterhouseCoopers, and I think Mark can actually comment on this a little bit, we had a discussion about this -- did an analysis about what the personalized medicine field is and whether or not you believe it's as broad as it is, but they estimate that from 2009 to 2015, we are going to see maybe upwards of 200 billion in growth in the personalized medicine market.

Now, at the core, we're talking about the diagnostics and the new drugs being developed, but as it broads out -- broadens out, it's looking at electronic health records; it's looking at clinical decision tools. In fact, they're moving out into complementary alternative medicine, health clubs, those areas, because the consumer market will more likely be drawn to that.

So, I'll move now a little bit into the regulation, a little slide of Calvin and Hobbs up here, you know, with Calvin arguing that we're getting in the way of scientific advances with all these stupid ethical questions we keep asking, and I think this is the big challenge that we face across the community. So, the second question we were asked was about the existing

institutional and regulatory framework set up to promote the development of personalized medicine.

I don't know that that's an easy question to answer, and I think that there are tremendous barriers in moving forward. From the science end, clinical research, we need to develop new models that can even identify subpopulations of patients who can respond to these drugs. From the FDA's end, they look at risk-benefit, the terms safe and effective, how are we going to model this for subgroups of patients to be safe and effective? Medicare and Medicaid, how are we going to pay for this? Intellectual property, what data can be shared, as we start moving into these different populations? And then ultimately, the market forces. What incentives do companies have to find out that their drug is not going to work in a percentage of patients out there? These are very challenging and very complex issues.

So, I can't really answer what the forces are, but I do want to talk about at least what some of the issues that I'm seeing are. First off, who owns the genome? This is something that's being debated -- well, will probably land at the Supreme Court here. A company called Myriad says that they own the breast cancer gene, BRCA I, and they charge $3,000 for a test. Well, the

cost of sequencing DNA will certain be about -- down to your genome down to about a thousand dollars in the next year -- few years. So, how are we going to make these two things fit? Will we actually deny someone to know their BRCA I status?

As a corollary, we're learning more about these diseases. We need to find ways to do drug combinations, develop them two at a time, three at a time. These are -- this is a huge regulatory issue, because we're also mixing different toxicities with that.

And then finally -- and this will be the last big issue -- but the reality is genomics is here. FDA is challenged right now with how to deal with things like direct-to-consumer genomics. Can we market these tests to individuals? We're seeing this incredible ability to sequence DNA. Should patients have a right to get their genome? As a corollary, should patients have a right to share that genome with researchers out there?

We are going to have to wrestle these. I don't think anyone has the answers just yet, but the fact is the technology is here, and it's moving very -- moving much quicker than the policies are. So, I'll end with this slide. If you have a chance to read this book by Clayton Christenson, it's called The Innovator's

Prescription. It talks about personalized medicine and that we need disruptive innovations to move to this.

As he looked at some of these diseases, things like, in oncology, with cancer, we are right only 25 percent of the time, the first time we give you a drug. So, we need to find better ways to do that.

So, I'll close with that, and I'll turn it over to Mark -- oh, to Clark. Clark. Thank you.

(Applause.)

DR. NARDINELLI: Thank you. I just want to say it's always a pleasure for me to come to the Federal Trade Commission, where I can be surrounded by economists rather than by physicians, not that there's anything wrong with physicians. It's just a nice change.

I'm going to cover much of the same ground from a slightly different perspective, from the perspective of an economist at the regulatory agency. So, let's see if this works. Yes. Okay.

As the previous speaker said, there can be broad and narrow definitions of personalized medicine. I will start by saying all medicine is personalized. That's what any physician would tell you. They get as much specific information about each patient as they can before treatment. So, what is new, I think, is the use

of pharmacogenomics, this really powerful new sort of information that can be used for informing personalized medicine.

Now, what is that used for? Well, screening for adverse events. We -- there are certain genetic markers who say who might react unfavorably to a particular medicine. The example I have worked with is warfarin. We see who is likely to suffer severe bleeding events with warfarin. Selecting dosing, dosing is trial and error. If we find correlations with particular genetic mutations or genetic dispositions, then the range -- dosing is still going to always be trial and error, but the range of the trial can be narrowed, at least in the early going.

And predicted biomarkers, which is really what I want to talk about a lot today, and this is kind of my narrow definition -- maybe it's not a definition so much of personalized medicine, but it's a definition of what personalized medicine can do and what its -- I think its most important contribution can be, and this is identifying responders and nonresponders. This was mentioned, it's part of personalized medicine, but this is really the big elephant that we're hunting, if it's still all right to hunt elephants.

I guess -- I don't know if I'm giving away a

state secret here, but for most patients, most drugs don't work most of the time, okay? That's the way it goes. So, this is what we're really dealing with. So, how do we find who it will work for, okay? And that's where I think predictive biomarkers, the companion tests, can come in.

Now, let me just give you a -- I actually brought some numbers. These are real numbers. There are currently -- and currently, I mean as of earlier this week -- 111 examples of pharmacogenomics information on prescription drug labels, okay? So, essentially, there are 111 bio -- pharmacoeconomic pieces of information out there on the label. This is the physician label, the one that goes in the box or is in the Physician's Desk Reference. I forgot to count up how many labels we have, how many unique labels we have. As most of you know or many of you know, the label on the generic is the -- the physician label on the generic is identical to the physician label on a branded drug.

I used to have this in my head, but it's in the neighborhood of -- it's several thousand, okay? So, in terms of actual FDA-verified personalized medicine information on pharmacogenomics, we're still really, really very early stage. This isn't a lot.

And then when we get to what I think is the real

important target, patient screening, identifying nonresponders from responders, there are only 18; although, as I say, 16 of the 18 were approved in the past decade. And of those, 17 of 18 are for cancer treatments. So, when we're talking about responders and nonresponders, it's only a very small exaggeration to say we're talking about cancer, okay? And, of course, the example that Adam gave was cancer, okay?

And most of these, of course, came after the fact. These were either tests for new drugs or -- I'm sorry, tests for old drugs that people thought might have uses or just ways to find something you could do for some of these really difficult cancers, cancers that were very difficult to treat.

Oh, and there is a mistake here. It says 2001, and that's a typo, it's 2011. This year, however, there were two new cancer drugs approved with companion diagnostic tests, and this is, I think, the future we should be looking to, the new drugs with actual diagnostic tests that are approved at the same time. They are bundled. One of these was for late-stage melanoma, and the other was for a rare form of late-stage lung cancer. It was a -- it's a form of lung cancer that is not associated with smoking. There is a small group, okay?

So, these are very new. Note, however, that they were both cancer tests. Some of you may be aware that it's been in the news recently that the New England Journal of Medicine has published a study with another companion test for -- and a drug to treat cystic fibrosis. It's a test that identifies people with a particular genetic variant. It affects about 4 percent. This would be another -- this would be a noncancer example of diagnostic screening through tests, but, again, this is -- this is very rare.

So, what we are looking at, then, despite all the promise, is still a very narrow use of these new methods, and it's largely in cancer, and it's typically for very small patient groups relative to the total, okay?

Okay. Well, the next step, as I've hinted at, is can we get predictive biomarkers for more noncancer treatment? The only actual one is a hematology drug for, again, a very rare type of anemia. Companion approvals, we have a draft -- the FDA has a draft guidance document that's still in process for procedures for companion approvals, the drug plus the diagnostic. We also have to worry -- and this, again, was also brought up by Adam -- about incentives. Do the incentives, on balance, work in favor or against

investment in predictive biomarkers and companion tests?

And particularly, I'm thinking from the point of view of the pharmaceutical company. Obviously, if you're a medical device company and you make tests, you have an incentive to develop a test if it might prove useful. But as was pointed out before, we're talking about something that will narrow the scope, narrow the patient population that a particular drug will be used for, and that runs very much against the blockbuster model of drug development, okay?

The ideal drug from the point of your view -- and from the purely financial point of view, of course -- of a drug company, making branded products, is a drug that everybody uses for the rest of their life, okay? That -- you know, and, of course, you can't quite hit that, but with -- with Lipitor, you can come close. And that's really the model.

But now we're saying, well, let's -- instead of trying for that, why don't you try to find drugs that will affect 2 percent of the population instead? And that's -- you know, in the simplest model -- and, of course, you can build in all kinds of complexities. That's what all you academics out there and you FTC researchers do. But in the simplest model, a companion test, from the point of view of big pharma, of a drug

company, reduces the size of your market, the potential size of your market, and we all know that the size of market is a very, very powerful predictor of innovation, okay?

So, the real question is, will policies and regulations need to change? You know, we have the Food, Drug and Cosmetic Act, with Hatch Waxman, with pediatric exclusivity, with tropical disease vouchers and other ins and outs. It has begun to take sort of a Byzantine look to it. So, I -- you know, who knows what the next -- the next iteration will be?

There might be -- well, there obviously are ways that the policy could deal with this, but I'm just an economist who analyzes policies. We don't -- we don't -- my job description isn't to suggest things. I'll leave all that to you. But I think it's clear that, you know, we -- there is -- there is a big gap here between, you know, the incentives and the policy and the promise of personalized medicine, at least as I've described it, which the real sticking point, which is can we separate responders from nonresponders? That's the single biggest gap in medical knowledge and in medical practice.

Okay. Thank you.

(Applause.)

MR. TRUSHEIM: So, I'm Mark Trusheim from MIT's Sloan School. I also worked in industry for a number of years, and you are going to hear a lot of the same things but hopefully some new evidence and some data as well. So, you've heard what is a stratified medicine, and it's this -- this is from Eli Lilly about five years ago, the right drug at the right time to the right patient, which Adam talked about as well.

At MIT, we had this sense of you have empirical medicines, those that you give like vaccines; you have the very individualized stem cell kind of vaccines, they take your own cells and grow them up and put them back into you, very individualized, truly personal medicines, if you will. We only have one approved example of that in the last year, called Provens. This was -- longer ago, this was Oncophage. It's only been approved in Russia so far. And then there's this middle area, which we call stratified medicines, which is really what Clark was talking about, combining diagnostics with the actual medical treatment.

And you've heard Adam talk a little bit here about major drugs are ineffective for many of the people. This is based on some work by Abrams and Silver, which also calls on the Spears work. It's remarkable. There is only about two or three papers

published in the world about what is the variance of patient response sort of at this meta-analysis, and is it a big opportunity or a small opportunity? It's just not very well studied, although anecdotally, people know it's very true. And if you're like me when I first looked at that, you're shocked at some of these numbers, right? And that's because the way we develop drugs today is we look at those who are treated versus those who are not treated or treated with a placebo, and we say, is the average difference one that's efficacious? Stratified medicine and personalized -- pardon me?

AUDIENCE SPEAKER: (Off mic.) What's pink and what's brown?

MR. TRUSHEIM: So, pink are people who don't respond. Brown are the people who do, all right? So, the pink numbers are those efficacy levels. These are what Abrams reported. I won't defend the numbers, but that's sort of what this -- this spread, which is up there, all right?

AUDIENCE SPEAKER: (Off mic.) I'm sorry. Which one is which again?

MR. TRUSHEIM: Sure. Oh, boy. I -- so, pink are people who die, all right, who don't respond, all right, the way to think about it. So, in hypertension drugs, about 10 to 30 percent of the people given them

don't see the therapeutic benefit that you would have expected if you read the drug label.

AUDIENCE SPEAKER: (Off mic.) So, those are done (inaudible).

MR. TRUSHEIM: Those are non -- they are ineffective, right? So, those are ineffective levels, all right? All right? Because most people think if the drug's been approved, it ought to be effective for you, right? And so, it's fairly surprising when you find out that many of these drugs actually are fairly ineffective, either for the initial therapeutic benefit or the longer term survival, all right?

So, what we've been thinking about is it's this -- in this hierarchy of a patient presents, a physician comes up with a differential diagnosis, you confirm that diagnosis, and then you wind up with what's the best treatment, and that's where this diagnostic comes in. So, it's not about finding out whether you're really at risk. This is our more narrow view of it, and it's very similar to what Clark was talking about but not completely identical. That's what we're calling stratified, and many people think that they will -- all drugs will -- or conditions will stratify over time, and you heard the Herceptin story.

We would argue you have to have both some

biological mechanism differences -- there has to actually be a difference. If the drug works for everyone, it doesn't matter, just give it to everyone. If it doesn't work for anyone, it's a lousy drug and shouldn't get approved, but if it -- if there are some that respond and some don't, which appear to be most drugs, all right, then you have this opportunity, all right?

You have to have multiple treatment options, right? If you only have one drug for your condition, even if the chances are one out of a thousand, chances are you are going to take it, right, particularly if you can't tell up front, right? So, stratified medicine doesn't really play if you're the only drug that's out there, and you need a biomarker. For statins, right, there is really very minimal difference in how, from a clinical standpoint, in both mechanism and whether it makes any clinical difference from one drug to another, despite all the marketing, perhaps, that's out there, right?

You have many treatment options, and you have some clinical biomarkers, but frankly, it's not worth the hassle of testing to figure out which statin might optimize your cholesterol level, all right? So, we think that will probably remain empirical, right? There

are those who disagree with us, all right, but that's our view.

And in SSRIs for depression, there is hugely different biological mechanisms going on there. There is multiple treatment options, many SSRIs. They all seem to have the same effect on their supposed target, and everybody responds incredibly differently, which means they're probably actually hitting some other target we don't understand biologically, all right? Unfortunately, we have no biomarker, all right? We have no way to tell other than you take the drug and see if you're less depressed in a month, right, as to whether the drug's going to work for you. That's a market that would stratify nearly instantly if someone had a good marker going forward.

You heard Clark give you some basic numbers, right, about how many drugs are out there. This is very impressive. If you see the growth rate, it's unimpressive. This were both injectables and orals. We did some other analysis with IMS, and it was convenient to break it out that way. You add those two things together, comes up to about $20 billion. The growth rates look really impressive until you realize the total drug market is $650 billion, right, and this is less than 3 percent, all right, of those drugs.

And what's also interesting, as we looked at that data with IMS, is while the U.S. initially led in this field and in its usage, in the last few years, actually, Europe and Japan are using personalized medicines more intensively. This is their usage of personalized medicine per thousand capita, all right? Their intensity is actually increasing over ours, right? We're actually declining in this space. I'll leave it up to other people to interpret whether that's because we overused them to start with -- thank you, Laura -- or whether Europe is now catching up to our bad practices and we're learning how to be more cautious about it, right? But nonetheless, that's sort of the story.

So, we did some work in association with one of Clark's colleagues at the FDA and -- at the FDA and an industry consortium and some other academics and IMS and some others where actually we tried to understand what was the complexity and the incentives, and that just got published this week for anyone who's interested, and Nature Reviews/Drug Discovery. The team here is the number of the companies that are up there, but more importantly, we took a very broad view, from R&D, through regulatory, all the way up through commercial incentives, and we tried to quantify this going forward as to what was up.

We linked five different models. Three models were clinical trial simulation tools for what were the size and the types of trials that needed to be done, and we had two different economic models. The MIT model was more of a deterministic model, and the IMS was a Monte Carlo simulation tool. We benchmarked them against each other to give them the same inputs, to come out with the same outpatients, so we were pretty comfortable with all that. That wasn't clear going in, by the way, right? I know that one of the FDA's interests was to whether different analysis techniques led to different answers. We've thought so far it didn't seem to.

We looked at the additional what's called all-comers approach and three different kinds of stratification; one called a rescue, you wait until it doesn't work in the all-comers and you try to find a subpopulation; one where you do dual development, you think you have a great biomarker going in, but you still do the all-comers just in case the biomarker doesn't work, all right, or it's economically not what you'd want to do it, you've got enough impact on the all-comers that you'd still want to get approved that way; and one where you focus only on the biomarker subpopulation, which is what the last two that Clark talked about that just got approved did. They looked

only at those that had the particular mutation in their clinical trials. They didn't look at all lung cancer patients, only the ones that were already targeted, for example, with this mutation.

We looked at three case studies, the granddaddy one that we talked about before of Herceptin. We looked at another one called Vectibix, and we looked at an Alzheimer's drug that's in development. That would be the blockbuster, right, because there was this debate that Clark would talk about, would it work in a blockbuster world or not? The chart, which is unreadable, has (inaudible) Alzheimer's drug at the top. The X axis scale is percent improvement of stratified versus the traditional drug approach, and those are 100, 200, 300, 400, and 500 percent improvements. These are not, like, marginal 5 percent kinds of deals.

And in Alzheimer's, we found a benefit of five times, right, that the expected net present value in Alzheimer's for a stratified approach, looking at the APO E4 negative gene type population was five times that of going without that, all right? So, it can work very dramatically in a blockbuster market, it can work dramatically outside of oncology, and that is, indeed, the clinical trial design which is being pursued right now by those developers.

The other thing you'll notice is not all cancers are the same, all right? Panatime maybe there in the middle didn't work out so well, all right? And I won't bore you with all the details, you can read the paper and ask me afterwards if you're interested in that.

But moving on to some of the policy questions, all right, that were important, also in that paper, what we did was we looked at about nine of the factors. We took a billion dollar NPV oncology drug, looked a lot like Herceptin, unsurprisingly, but we fudged it up a bit from that to make it a bit more generic, and we took nine factors -- again, a little hard to read on this -- but it was everything from development, time, and cost, to cost of capital, to what was the pricing and the shares that one would get through that, and I could turn and we could turn billion dollar drugs into a $250 million loser by only changing each of those factors by 25 percent in the negative, all right, because these things all compound throughout the development process, all right?

So, this -- but what this leads to, and you'll see in a moment, is individual policy and external players optimizing in their factor area, doing very sensible things, can do things that actually confound and lead to what we jokingly refer to as pharmageddon,

all right, through all of this.

There is a Nirvana case, all right, where if you simply increase each one of those factors by 25 percent, I can turn the billion dollar drug into a $10 billion winner from the producer's standpoint. So, if you're thinking about incentives, all right, for what might make sense or not to do with this, and this has to do with everything from development times again to cost of capital, all right, or to the availability of capital, all things which both government and private sector and regulators and payers have a great deal of control over.

Being from MIT, we aren't comfortable with just doing that. We love hundreds of thousands of numbers. So, we ran the Monte Carlo approach where we did the 500,000 potential solutions, right? If you do 12 factors instead of the nine and you do a high, medium, and low for each, and you do the combinatorials, that gives you 531,144 different combinations. You plot those out. We turned this into expected net present value, taking in probability of technical and regulatory success, because that was one of the additional three factors that we used. And over half of the scenarios turned relatively uneconomic.

The expected net present value was about $150 million on this. We took a cut-off of about 100 million

as being marginally interesting, right? That's the yellow bars, stop at the hundred million dollar level. They actually go negative and turn red. And there's a very long tail of where it could be very exciting and profitable. And this was not that we were changing the science underneath, this was not that we were changing anything to do with the actual drug and diagnostic combination. This all had to do with pricing, with how long did it take to develop the drug from a clinical trial standpoint, what did the payers' response to all this look like, what was your cost of capital going through this? And you wind up with a world that can be very exciting or very discouraging.

This slide has the list of what those -- all of those 12 factors are, and some -- how many of the people respond to the drug, policy has very little impact on that, although more than you might think, because what regulators decide is efficacious, right, and what that cut-off value is for did someone respond or not respond is actually somewhat judgmental, right, as to where you'd like to place that. There is also, does it have to be a superior response or simply similar to what other drugs have already shown? That's been a place of great policy debate as well, all right?

Things like development time and trial size,

regulatory and payer requirements for how much evidence they require to decide to reimburse and approve drugs is actually directly tied to how long it takes to develop a drug and how large those trials need to be and what style of trials that they need to be, right? So, those things, while FDA will, of course, say they don't design the trials, right, and payers will say they don't tell drug developers what evidence they have to bring forward, in essence, they do, all right, because they say that unless you meet certain thresholds of evidence, and we require it in certain forms and in certain cohorts, we won't accept the drug going forward. So, this is an interplay between both the developers and the policy-makers through all this.

And increasing pressures on economic incentives, if you go around that feasible space, all right, of what's possible, on the regulatory side, in the past two years, we've seen the FDA decide that lab-developed tests are now subject to their review. That had been a place of great innovation. That is now becoming a place of less innovation or certainly less investment going forward, because of uncertainties there. There has been multi-variant test guidance that's been put out about if you use some of these gene expression patterns and you want to use more than one at a time, the FDA has decided

that the standards for those kind of tests are going up and are much higher; therefore, making costs of development longer and higher and less incented.

And the FDA does not accept any retrospective data.

Think of that as if you were Toyota or GM and you were doing quality control work, right? You all look at what the car did and what your historic failure rates were, in doing all that kind of Six Sigma work, that's not allowed in healthcare, right? You can't go back and do that kind of retrospective Six Sigma work and ask for an FDA label change, except in the case of safety, which is how the Vectibix K R A S change was made. It was adopted by all the clinicians, it was adopted in Europe, it would not fit under FDA regulations, as they were now, but they will accept it as a safety change, and that's how they snuck in.

Reimbursement on both drugs and diagnostics are both asymmetric. If you stratify after the fact and, therefore, add more value to a smaller set of patients, you are basically not allowed to increase your price point, okay? So, if I was only -- if I was only helping half the people and I had a certain efficacy level, say one-year life extension, right, then I can find the half that really responded for two-year life extension. I

can't double my price for those people to get to the same marketplace. I'm stuck at the original price that I was at.

There are some provider adoption issues and there are some exclusivity challenges as well that are there. There are many incentives, other than price, that one can use. We have some tried and true ones. They have not been applied yet to personalized medicine. And there's some new tools, again, around some adaptive licensing, some different kinds of trial designs that could be possible that people are beginning to move towards.

But in general it's been a challenging world, and that's just for the new product development, much less for those 1200 or so drugs already approved, getting any kind of stratification on those. So, that's sort of the overview. Let's see where we go from there.

(Applause.)

MR. TRUSHEIM: Did you want us to come up here, Chris?

DR. GARMON: Yes. If the panelists would come on up, I have a few questions, of course, if people in the audience, if you have questions of the panelists --

AUDIENCE SPEAKER: Could I ask a clarifying question? Why can't you change your price? Why

couldn't a drug maker change his price?

MR. TRUSHEIM: They always attempt, but the payers nearly universally deny that, and because the -- the drug makers generally in society can't withhold their product from a marketplace without ethical concerns, it's a asymmetric bargaining situation.

AUDIENCE SPEAKER: You mean the government, like in Medicare, the Medicare system wouldn't take a price increase, but would private insurers?

MR. TRUSHEIM: Private insurers, in general, will accept a few percentage point increases, right, from year to year, but going in with massive price changes has been universally unsuccessful thus far.

Yes?

AUDIENCE SPEAKER: So, I'm a little unclear on the model you have in mind in this discussion, because I thought we were dealing not so much with a product that had been approved as safe and efficacious for everybody and then you discover that it's only for some, but more interested in the possibility that a product that didn't pass the statistical standards for safe and efficacious for everyone could be safe and efficacious for some.

And that raises, to my mind, the question of incentives and statistical testing that I think Clark raised, but, you know, if -- if we're dealing with a

product that will help some people and not others, if you're looking at the overall question of does it help people, you may well find a negative answer if you look only at the whole population, but then you can -- could, in principle, do a bunch of data mining and find some marker that might be correlated with the people that it happens to help in the trial sample, and the question is, when should you accept that as evidence for a targeted efficacy finding?

And obviously that's going to change the way you have to think about false positives and false negatives, and it will change the incentive to get into this in the first place.

DR. CLARK: If I can just comment on that, I think that that is one of the big challenges that we have Avastin is -- you know, we had a big controversy this summer. Some women benefited. On the whole -- and this is for metastatic breast cancer. On the whole, it was accepted for accelerated approval, and FDA said, on the whole, the risk-benefit does not pan out.

From the patient community, there was a big outcry of now you're denying us a medication, and the FDA said, we don't have the markers to predict who it's going to work in and who it's not going to.

Provange, similarly -- so, this was the prostate

cancer new type of therapeutic that uses the individual cells had -- the cost was around $90,000 for treatment, and the benefit ranged from four months of increased life to about 2 1/2 years, so broad range. And, again, we don't know who's going to be on that 2 1/2 years, and we don't know who's going to be on the four-month, but when you look at $90,000 for four months, how do we pay for some of these? And it's an incredible challenge, I think.

But getting back to your point and that Mark brought up, working with the FDA to do retrospective analyses on some of these samples, I think, is something we need to find better ways to do.

DR. NARDINELLI: Let me add that the example you gave -- there are examples of this. You know, this is a case where the incentives between the -- to develop the diagnostic and to find the biomarker merge with those of -- of the pharmaceutical company. They've got a drug that isn't for everybody. Maybe it is for somebody. So, if that's the case, it's kind of a -- but it's -- I think it's an important development, it's going on, but it's still kind of a consolation prize for also-rans, and -- from the point of view of the pharmaceutical company. Obviously, from the point of view of patients, it's a very good thing.

AUDIENCE SPEAKER: (Off mic.) But isn't it true that most drugs that look promising turn out not to show that they work for everyone? So, the consolation prize is more common, in some sense, than the (inaudible).

DR. NARDINELLI: Well, working for everyone isn't the -- you know, the criteria for a drug to be given to everyone. So -- but -- so, it's more a matter of safety.

MR. TRUSHEIM: So, the first few examples may have been what I call these retrospectives rescues, right, which is sort of your model. The industry and many people are trying to move prospective, right, when the two examples that were approved this year were that sort of prospective approach. And then there's also the drugs that have been approved historically, right, where we saw some data today that says they don't work uniformly, and did we want to have incentives for either diagnostics or the drug companies to go back and look at older drugs to help better target those so we treat people better and spend less money on treatments that don't work?

AUDIENCE SPEAKER: So, I just wanted to follow up on what was -- on both what Joe was saying and the previous discussions about pricing, which is that, you know, how many things fall into Joe's situation, where

there's actually a big incentive for the drug maker to develop the diagnostic and how many fall into the other situation, to depend crucially on the pricing, because if you have freedom to choose pricing, there is actually a very clean economic theory that Mark has contributed to and that Lewis and Sappington and Johnson and Myvin (phonetic).

There's a whole series of papers about when you do or do not want the consumers to be informed, but when you can't increase your price when you inform them, because that's the standard thing that you do, that dramatically reduces the set of cases in which you have an incentive to make them informed.

And so, I think both relating these sorts of thinkings to that literature and thinking about how sort of the price cap affects that would be a nice economic research agenda.

DR. GARMON: On that point, I'm wondering whether the -- since a lot of these therapies involve a diagnostic test to find the biomarker, whether the pricing of that test and if that could give the incentives, whether you would need to have that also patented by the company that's marketing the drug, if that's a solution to the price problem.

MR. TRUSHEIM: So, did you want to comment on

that first? Okay.

So, currently, in the diagnostics market, it's a cost-plus pricing schedule that CMS has set and that most private insurers follow. So, if you develop a new genetic test, you know pretty much ahead of time what your pricing will be. It's about a hundred bucks, all right? It doesn't matter whether it's a test that adds a tremendous amount of value to society or adds generally no value. The price for running the test is a hundred dollars. If it's a protein immuno assay, it's $36.83, all right? And it is fixed, regardless of the value that is provided.

Those firms that have chosen to go forward, like Myriad with the BRCA I and II tests, and say, well, we will not sell it at that price. We have the IP. You can only get it from us, all right? We will sue any academic lab who attempts to run the test, even though they could, because that's what our patent system allows, all right? And they charge two or three thousand dollars for value pricing to help people avoid a 50 to 100,000 dollar breast cancer incidence, and they are uniformly denounced as being not just gouging on prices, but it becomes very much an ethical -- they are somehow evil for pricing a diagnostic test above what the standard schedule might otherwise allow.

So, there's a very strange dynamic in the healthcare markets, that diagnostics are supposed to be cheap, if not free, regardless of what the value is that they deliver to the overall system, but other parts, whether it's surgeons with surgeries or drug companies with drugs, can routinely get tens of thousands of dollars, and that's perfectly fine, right? So, there is some weird asymmetry that I would love to hear economists talk about that's going on in this marketplace.

DR. NARDINELLI: Yeah, and it's -- also, the diagnostic industry, as part of the medical device industry, has historically been very, very different from big pharma. It's relatively small firms, a lot of entry and exit. It's engineers, not physicians, and there's still -- despite the -- you know, the relatively low pricing, there is a lot of activity. It's a very -- you know, a lot of entry, a lot of things going on.

So, there -- we are talking about maybe a very -- some of the suggestions really would imply a very different approach to the industrial organization of this market.

DR. GARMON: On that point, do you think that the new payment reforms, with the healthcare reform and other payment reforms that are being used by health

insurance companies, bundled payments, accountable care organizations, will those things help in the development of personalized medicine or harm the development?

DR. CLARK: I don't know that I can answer -- not being an economist, I don't know that I can answer that. I think what we're seeing not only with the ACOs, the integration of health information technology, the adoption of electronic medical records, that can assist in clinical decision-making, is going to change the model. And what I mean by that is getting -- as we start to learn more about personalized medicine approaches, using genetics, getting better tools to doctors to help make some of those decisions. It's going to need to be balanced by some type of reimbursement. If you go to a surgeon, their recommendation is going to be surgery for a procedure. If you go to, you know, a medical oncologist, it might be a very different type of treatment. It's going to need to interact with what the patient's needs and expectations are, as well as what the tests are telling us.

DR. NARDINELLI: Well, again, the FDA doesn't -- we don't do prices. So -- but I would think that any sort of new pricing options could only help.

MR. TRUSHEIM: Yeah. If it can overcome -- you

know, certainly the incentives will be to practice more integrated medicine, right, which is a good thing. The question is, in my mind, whether it can also incent some of the behaviors that were existent decades ago. People may not remember, but the first mammogram screening clinical trial was sponsored not by a drug company or an instrument maker, but was sponsored by a health insurance company. New York Life & Health Company sponsored the first clinical trial for mammogram screening for breast cancer, all right? That part of the marketplace has ceased investing, all right, in clinical studies that would seem to be greatly in their interest to understand which drugs worked for which people, which surgeries were most efficacious. That would seem to be a fantastic opportunity for them to lower their costs and their medical loss ratios, and they don't invest in that kind of R&D work.

I'm sure there are questions about how long they have lives under control and whether they -- there's a lot of spillover effects and free-rider problems, et cetera, with that kind of value capture, but some of them are large enough that you would think that the free-rider issue would be, who cares? I'm still going to save a huge amount of money if I do it. But that hasn't been their culture, and we don't see any evidence

that that's going to change.

AUDIENCE SPEAKER: (Off mic.) (Inaudible.) I mean, if you are saving a lot of costs that are post-65, the incentives have gotten weaker, not stronger.

MR. TRUSHEIM: Yes, that is absolutely correct. There is -- but we don't see any evidence in the younger diseases that they're investing either, right? So, it -- and with CMS and NIH and others, we haven't seen that kind of attention that that would be a great benefit for the government to be investing in those kinds of trials to lower their costs. So, there's lots of opportunities for creativity as to whose interest it is, but we haven't been able to overcome whatever institutional barriers and culture problems prevent that.

AUDIENCE SPEAKER: So, I can see an argument for CMS, but if it's going to -- if all insurers were then to institute this test or this procedure, then that's just going to -- if we think these markets are at all competitive, drive down prices. So, the person who did the study would reap very little of the benefit.

MR. TRUSHEIM: That is true to an extent, you are absolutely right, but the pace of change and adoption in medicine, the dissemination of new technologies is measured anywhere from five years to 30

years, right? So, if I had some controlling and good evidence-based medicine and driving down those costs, that would be great. And, in addition, if you think that, in general, that the industry is having pricing problems, all right, with private payers not being able to sell to employers, that they can afford to buy their products, right, that the health insurance had gotten too expensive, a general lowering of the -- of their costs would expand the marketplace tremendously for them, and, yes, they would still be competing, right, with others, but their marketplace would expand, and they would also blunt some of the move that seems to be going towards greater and greater government coverage for larger swaths of the population.

But that kind of argument, most of the insurance executives agree completely with you, right, and that that's the way the market is currently structured, right? They don't see any advantage to that, and they don't see any advantage in banding together to do anything about it, either, if they believed the larger market expansion story I just wove.

AUDIENCE SPEAKER: I think you are missing an obvious reason why they don't want to do the studies. If they tell the statements we're not covering this treatment because the study shows it doesn't work, the

patients aren't going to be grateful. They are going to be angry. So, why?

MR. TRUSHEIM: Well, I don't know too many patients that don't -- wouldn't like to know that the treatment's not going to work for them, all right? There are the two examples we've seen in the last year with the mammogram study and now the prostate surgery ones, which had a huge, oftentimes government-funded PR effort behind them, right, to convince consumers that these were greatly in their benefit to do. And now, when the evidence has come out that these screenings, perhaps, are not justified, from a survival standpoint, that there's backlash to that. I think those studies, by the way, may have been flawed from a consumer standpoint. They only looked at survival. They didn't look at quality of life during that time and a number of other things, and that's another huge challenge in healthcare, is that consumer and patient preferences generally are not factored in at all, right? It's the old Ford Model T. You can have any color you want, as long as it's black, and I only make one version of the car.

DR. CLARK: And I was in the middle of all of that particularly when I was at LiveStrong, particularly with what was going on with the mammography issue, and

one of the big challenges is also from the public's end, a misunderstanding of what population-based approaches are versus individual-based approaches. The -- those who may have had a family history of breast cancer or of prostate cancer. The perception was we're going to deny this screening for everyone, versus a body that was looking at a population number with end points that are population health, not individual health.

DR. ADAMS: I used to play football against Adam, so I am going to tackle him. One of the things with personalized medicine, you are going to make the quantity smaller, the only way to get the incentives to develop the drugs is to make the prices a lot higher, and if we're talking about 200, 300,000 dollars per person, per drug, organizations like your former organization have come out against such pricing. You know, how does it all fit together?

DR. CLARK: Well, I think it's a huge challenge. What you've presented, I think, is the current model for drug design that was built on the blockbuster drug model. The other way would be to try to reduce costs in the front end and have different companies looking at how do we redesign this model so that it's cheaper to make drugs more efficient? Right now, it's about a billion dollars to develop a drug, takes about 14 years.

From the patient's end, they're well aware that costs too much and it's taking too long. Are there different ways to look at it?

I'm a big proponent of health IT for just this reason, particularly with the integration of genomics and genetics. If we can start to identify the potential responders based on whether it's screening in the HR or using genetics, get them into trials for some of these new drugs, hopefully the cost would be reduced and the time would be reduced. There might be a different economic model. I'm not an expert, so I can only argue that to -- you know, so far, but I would hope that that would be a way to go about -- about this.

MR. TRUSHEIM: Yeah. And to back that up, in our 2007 paper, we ran such a model, right, and you can change the world from needing a billion dollar blockbuster to you can easily break even or -- from an economic equivalency standpoint be just as effective with return on investment of a $200 million product, all right, if you sped up the process, right, and also lowered the cost. And we also have patent expiry issues which are -- oftentimes, you only have ten to 12 years, maybe only seven years, to sell your product in the marketplace. That's, again, a legal structure we've put in place for drugs on the basis of the patent system.

For media, all right, in terms of music and others, that is a 70-year window, right, to 110-year window, right? A very different time frame. So, if you only have ten years to make all your money back, it's very different than if you have half a century to make your money back.

DR. GARMON: So, on that point, why doesn't the Orphan Drug Act and/or the new biosimilars legislation, which has a very long data exclusivity, in my opinion -- a very long data exclusivity --

MR. TRUSHEIM: Others who disagree, but yes.

DR. GARMON: -- why wouldn't that solve the problem? Why don't those pieces of legislation solve the problem?

DR. NARDINELLI: Well, to some extent, we're not solving the problems, but many of the drugs we have talked about or have been described have come under the Orphan Drug, so that the -- one answer is that the Orphan Drug is being used in these cases, okay? But there's a wide range of things that it's not fitting. So, it's -- you know --

DR. GARMON: And why is that? Is that because of the diagnostic test? Is it the Orphan Drug Act just applies to the pharmaceutical itself and there is no way to (inaudible)?

DR. NARDINELLI: Yes. Actually, we need a

lawyer here to really explain the Orphan Drug thing. But yeah, there are -- it applies to some things and not others. On the other hand, there is, under the Orphan Drug Act, a little known provision that it can apply to a drug that would otherwise not be financially viable. It's often thought of as being only for small populations. So, some of this can be done under the Orphan Drug Act.

MR. TRUSHEIM: The challenge with the Orphan Drug Act, aside from the escape clause that Clark just mentioned, you have to be classified as an orphan disease, all right? It's often ambiguous to the FDA as to how to interpret what is a disease, all right? So, is a certain genetic mutation in a multi-million patient population, is that a new disease or is that simply a subpopulation, like pregnant women and African-American, right?

DR. CLARK: And I think this is actually where a lot of the patient communities are going, particularly the cancer community, is we're learning -- and we're now talking about there are hundreds of types of cancers based on the genetics. So, the Orphan Drug Act I think applies to $200,000 or 200,000 incidences or less, and I think arguments can be made that many of these cancers are actually that. It's not just breast cancer. I may

have triple node negative breast cancer, which is very different than other types of breast cancers, and these are some -- again, from the disruptive innovation model, I think some of the things that are going to push that limit. When the regulatory environment is restrictive to some of this, how do you work outside the regulatory environment to push a personalized approach?

AUDIENCE SPEAKER: So, we talked a bit about the incentives for various possible actors to develop basically use-restricting tests in the U.S. market for a drug for which the but-for world is it's taken by everyone. Maybe the answer is there aren't enough people with enough incentives to do that in the U.S. market, but presumably, a drug that is on the market and being used for everyone with a disease, some of the single-player -- single-payer systems in other countries might have an incentive to do that part of the work.

And so, are there unnecessary or -- or possibly removable obstacles to that kind of international decentralization of who does what?

MR. TRUSHEIM: So, most drug developers think globally today and look for global markets, all right? The FDA and their European counterpart, the EMA, have done remarkable work trying to harmonize their processes and their criteria. So, while it's not a uniform data

package to each, it's reasonably similar, all right, to each other.

Asian countries are moving much in that direction at times as well, although unevenly and differently. Your single-payer argument, right, in Europe would seem to have a lot of merit. The challenge institutionally appears to be that those who are paying for treatment are like CMS, right, and those who are funding any potential drug development look like NIH, and they're very different arms of the governments over there, and they have the same kind of payer challenges, even in the single-payment systems, that they have not structured their funding such that they have, it appears, any surplus in those CMS-like payer organizations to invest in any other kinds of research. They just have set themselves up that whatever the medical costs are is what they pay, and they don't really have a 5 percent or a 10 percent R&D budget, in essence, coming through those single-payer systems. So, there seem to be some very institutional, structural challenges that I've observed. Those of you who know more about it can probably say much more insightful things, but that's a layman's observation from being in -- in the industry.

DR. ADAMS: Okay. Why don't we leave it there

and thank Chris, Adam, Clark, and Mark for a very interesting discussion.

(Applause.)

DR. ADAMS: At the start of this session, Joe raised a challenge that there are a lot of very intelligent academics out there not working on very interesting questions. I think if you looked at our agenda, you found a group of academics who are working on a -- some very interesting and very important questions for the agency. So, I want to thank everybody that was involved in this conference, the scientific committee, Mark, Nancy -- and I'm going to forget who the scientific committee was -- Aviv, right, who could forget Aviv? -- and who was the last one -- David, the Northwestern people anyway.

Also, I want to give a big thank you to Laura Kmitch, who is one of the RAs at the Commission and does a fantastic job putting this conference together. So, again, I think a big round of applause for everybody for a great conference.

(Applause.)

(Whereupon, at 12:57 p.m., the conference was concluded.)

C E R T I F I C A T I O N O F R E P O R T E R

DOCKET/FILE NUMBER: P085800

CASE TITLE: MICROECONOMICS CONFERENCE

DATE: NOVEMBER 4, 2011

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the notes taken by me at the hearing on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: 11/23/11

ROBIN BOGGESS, RMR

C E R T I F I C A T I O N O F P R O O F R E A D E R

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

SUSANNE BERGLING