FEDERAL TRADE COMMISSION

I N D E X

SECOND ANNUAL

FEDERAL TRADE COMMISSION

& NORTHWESTERN UNIVERSITY

MICROECONOMICS CONFERENCE

FRIDAY, NOVEMBER 20, 2009

FEDERAL TRADE COMMISSION

601 NEW JERSEY AVENUE, N.W.

WASHINGTON, D.C.

P R O C E E D I N G S

- - - - -

DR. ROTHSTEIN: All right, good morning, everyone. Two quick announcements. Please note the evaluation forms out front. We would certainly appreciate it if you would take a moment to fill out the evaluations for the sessions that you've attended. You can return them right there, there's a basket for them.

Also, a couple of papers with notations on them were collected yesterday, so we have them. If you would like your insights back or you wanted to copyright them, they're not yet copyrighted, you can collect those, too, up front. That's the extent of our lost and found.

It's a pleasure to introduce our next keynote speaker, Kyle Bagwell. I introduced him yesterday, so I assume even after dinner the information hasn't changed, unless you sent an email you shouldn't have sent. He's still the Donald L. Lucas Professor of Economics at Stanford, and Kyle has written on a wide range of topics. He's written about the general agreement on tariffs and trade, the precursor to the World Trade Organization, the value of price as a signal of quality, the economics of collusion, the economics of advertising. He's a fellow of the Econometric Society, which is really quite an accomplishment. He is the

editor of numerous journals and is a former editor of the Rand Journal of Economics. He has taught at Northwestern University and Columbia University, as many of you know, and he's co-author of various books, but in particular The Economics of the World Trading System with Bob Staiger. It's a pleasure to have him here this morning.

DR. BAGWELL: Thank you very much, Paul. It's very nice to be here and thanks for your kind introduction.

I thought today I would talk about some work that I have done, mostly with Susan Athey and also with Chris Sanchirico on collusion and private information.

To set up the discussion, I thought I would begin by highlighting some of the ideas that underlie the basic model of collusion, which would be a repeated game with full information. If you want, think of a repeated Bertrand game where firms are interacting over an infinite horizon, choosing their prices day after day. The basic incentive constraint that arises in that game is that there's always a short-term incentive to undercut your co-conspirator, and you balance that against the long-term fear of triggering a price war.

Out of that perspective comes a lot of useful comparative statics, things like it's easier to collude

when there are fewer firms, when firms are more patient, which might correspond to a rapidly growing market, and various other results with which we're all familiar.

One highlight in subsequent research on collusion is the Rotemberg-Saloner literature, which talks about a setting where demand bounces around day to day in a publicly observable way. In that model, if quantity demanded is very high on a given day, then it's going to be very tempting to cheat. The best you are going to be able to do is to collude at a low price, assuming demand shocks are distributed independently through time.

A second important literature is associated with the work by Stigler, Green and Porter, and Abreu, Pearce, Stacchetti. This work relaxes the assumption of perfect observeability and assumes that actions are hidden. So, firms choose output and their output choices are not observed. Firm one doesn't see the output choice of firm two, there's demand shock, out pops a public price, and everybody sees the price. Then if the price is low they may say, well, gee, I wonder if the other guy cheated and produced too much, or if instead the demand shock was small. As a consequence of this inference problem, this literature predicts that we should sometimes see periodic price wars. So, firms are

going to episodically go through price wars. Some recent work by Joe Harrington and Andy Skrzypacz on private monitoring also fits within the category.

Our work comes in within the context of private information. Here, all of the actions are going to be publicly observed: Firms are going to choose prices and everyone sees the prices, but what's going to be different is that firms have private information about their costs. You could imagine other things that they might have private information about, but we're going to assume that they're all privately informed about their respective costs. The model is thus a repeated game with adverse selection or hidden types.

There is an old literature, somewhat old, from the eighties, I'm more hesitant to say this is an old literature, now that I'm older myself, but in any case there is an older literature from the eighties on static mechanism design work that looks at a related problem. This literature says that if firms could write a contract and they have private information about their costs, then they might write a contract that says something like, if I've got low costs and you've got high, let me do the producing, but to make sure that you can trust me when I say I've got low costs, I promise to give you some money, a side payment, that I would only

be willing to pay you if I truly had low cost.

That would be a way for the firms to try to find an incentive compatible means of revealing their costs and colluding, but that literature didn't address how do the firms enforce collusion if they can't get the government to enforce their contract in a court of law. Also it utilizes the monetary side payments.

So, we might think, well, what would happen if firms have private information and they're trying to collude, but they don't have side payments because maybe they're worried that the side payments would leave a paper trail that would alert the antitrust officials, and they also can't get the government to help them enforce contracts. Their agreement must then be self-enforcing.

That's where we come in, and I'm going to give an overview of three papers in 20 minutes. That's already sounding implausible, but I'll touch lightly on all three.

All three papers are joint with Susan Athey. The first paper I'll talk about is also joint with Chris Sanchirico.

So, the motivation of this work, some of which I've touched on, is just to add some greater realism, by introducing private information, and to look for new

predictions. When we look and see the conduct that the collusive firms would engage in in this setting, that the finding might provide some guidance about the sort of screening that one might want to use. The findings might also provide new predictions about the relationship between collusion and welfare, and they might also give us some insight about the role of communication in collusion. Communication, of course, is a critical component of the legal approach towards collusion. Have firms reached an agreement on price is a critical question, but getting communication to do something in the theoretical treatments of collusion, is tricky. So we're going to follow Bonnie Raitt's advice and give them something to talk about. In our models, they may talk about their private information about costs.

Methodologically, this work is kind of fun because it draws on the repeated games and mechanism of design literatures, and in that context you might just imagine that continuation values will play the role of transfers or side payments from the earlier literature that I mentioned a moment ago.

So, I'll start with what was the first paper. It wasn't the first to get published because it took a little while in the revision process, but the first

paper that we wrote was by Susan Athey, Chris Sanchirico and me, and it was published in the Review of Economic Studies in 2004. It's entitled "Collusion and Price Rigidity." We analyze the repeated Bertrand game that I was talking about. The firms choose prices, the prices are publicly observed, and the firms have private cost shocks. The cost shocks are independently and identically distributed across the firms and over time, so every day each firm gets a new cost shock. We considered both inelastic demand and downward sloping demands, but our results are strongest for inelastic demands, so what I tell you I will just highlight that case here.

The solution concept that we used here is worth noting. We focus on something called the strongly symmetric equilibrium solution concept, and what that means is that starting at any given period, the firms perceive the same value to the relationship. You don't start in a given period and say starting from here firm one is going to do really well and firm two is going to do poorly.

Rather, starting at any given point in the game, the equilibrium concept requires that all firms perceive the future as being of the same value in expectation.

So, you can have price wars in this setting

under the solution concept, but they would be industry price wars. What this is ruling out is that you play tomorrow in a way where firm one does really well and firm two does less well.

I will come back to this later in the other papers, but this first paper uses strongly symmetric equilibria. Wherever the firms go tomorrow, they go there together. They thus use symmetric strategies and have symmetric continuation values.

Transfers in that context, then, can be thought of as price wars. Let's suppose I claim that I have low cost, and we have some collusive scheme that says if some firm says it has low costs and wants to produce a lot of output, then we discipline that by sometimes going into a price war in the future.

You can have that in a solution concept, but from the cartel's point of view, that would be a transfer that's out of the cartel. The price war would be money that you're burning and sending off to consumers in essence. It wouldn't be a transfer from one firm to the other firm, as the symmetric solution concept doesn't allow that.

So, in this context, the goals of collusion are two: The first is the standard idea that the firms want to get together and raise the price. Second, they may

want to achieve productive efficiency, too. They may want to get together and say, who among us has the lowest cost and let's make sure to maximize industry profits that we assign production to the firm with the lowest cost.

But, of course, that's where the incentive compatibility question comes into play. If you're going to get more production from lower cost firms, then everybody is going to want to claim that they have lower cost, and so you are going to need some way to make sure they are honest. That's a problem that confronts the cartel. How to elicit honest communication about cost positions when cost positions are problematic?

Here's a picture that highlights some of the themes that are involved. So, Pn of theta is the Nash price as a function of cost. Think of theta as the cost you draw as a firm in this cartel. You draw some cost, and if it was a static noncooperative equilibrium, so that you weren't trying to collude at all, then higher cost firms would probably charge higher prices than lower cost firms. This Pn of theta function is upward sloping, capturing that idea.

This would be the pricing function that all firms would use in the noncooperative equilibria. So, let's evaluate it. Let's suppose that we're an emerging

cartel. We get together, look at the Nash pricing function that we've been using prior to our meeting, and say what can we improve? Can we do better?

We might look at this Pn of theta function, and say, well, one good thing about it is that we do get productive efficiency. Whoever among us has the lowest cost, that's going to be the firm that charges the lowest price, and that's going to be the firm that, therefore, makes the sales. The prices are going to ensure that, as long as the pricing function is upward sloping.

As cartel members, we might also say that the bad thing about this pricing function is that the prices are low. We're being very competitive with each other, trying to undercut each other, and so, from the cartel's point of view, low prices are the down side of the noncooperative equilibria.

The variable r that you see illustrated, think of that as being the consumer's inelastic demand, that's the reservation value for consumers. You can see that Pn of theta is potentially quite far below r, and the firms might look up longingly at r and say, wouldn't it be great if we could price at r, that would be so much better for our collusive profits.

So, the first thing they might try, if I can

point the arrow, is this upward sloping line, maintaining productive efficiency, but pushing the price higher and closer to r. That would be kind of an ideal situation for the cartel. They get basically the highest price that they could get, r, but add a little slope to make sure that whoever has the lowest cost makes the sale.

That would be nice if you could do it from a cartel's point of view, but the problem is that the higher cost firms are going to look at this and say, if I just cut my price a little bit, I increase my chance of winning a lot, and so why would I want to admit that I have high cost and charge a price up here. Better would be to cut my price down a little bit, so that I end up pricing as would a medium cost guy. Then I increase my chances of winning a lot, and if anyone ever challenges me about this, I'll say, oh, I had medium costs, that's why I did this.

So, this second candidate that I have presented for you, this somewhat flat upward sloping line close to r isn't incentive compatible. Firms aren't going to price off of this schedule honestly. What you might do, if you wanted to implement this, you might say, well, if we see these lower prices down here, then we're going to go to a price war in the future. That would be enough

to stop a high cost firm from dropping down. He would think, if I cut my price, I'll increase my profits today, but I might trigger a price war. I don't know, it's not really worth it. I think I'm going to stay put where I'm supposed to be. But that just highlights that if you're going to get productive efficiency as a cartel, then you're going to have to lower prices, like in the static Nash equilibrium, or maybe you get higher prices but you have to go through these price wars. So, it's not cheap.

Assigning production efficiently involves some incentive costs. So, what you might do instead, what you might do instead is pool. As a cartel, you might just say, I am going to give up on achieving production efficiency, I'll let everybody price at r. Then we get the high price goal of collusion, but we give up assigning production efficiently among our cartel members.

That's an option. You might think it's just too painful to try to achieve productive efficiency. Either we have low prices or we have price wars, so let's just give up on that and pool with everybody charging r. If our cartel works that way, then we split the market evenly. That would be an option. That would be another thing we could do.

So, there are all sorts of candidates that a cartel could consider. These are just three. Another option would be to try step functions. Other nonstationary equilibria, where you go in and out of price wars, could also be tried. There are a number of things that could be tried. The question is, what's the best way for this cartel to collude in this setting.

So, here's the result that we have: If the distribution function of types is log concave, which is not a free assumption, but probably not terribly restrictive, and if demand is sufficiently inelastic, which is the case I've been describing, and if firms are patient, not infinitely so, but sufficiently patient, then rigid pricing and stable market shares is the optimal strongly symmetric equilibria for colluding firms.

So, in other words, as a cartel, you just say the best thing we can do is to let everybody price at r and split the market evenly. We get the high price part of our collusive goal, and we give up on trying to achieve productive efficiency. We could achieve productive efficiency, but it would be too costly in terms of the incentive cost and low prices and price wars that would be required.

Under weaker conditions, you never want to use

price wars. For the cartel, price wars are wasteful and entail transfers out of the cartel, so the cartel generally doesn't want to use those.

Suppose you were trying to use some sort of screen and say, where should I look more closely if I suspect collusion? Then according to this line of thinking, markets in which the variance of prices is low, and price levels are high, assuming you have some estimate about what the costs are so that have some informed view on what a high price would mean, those might be markets where you would want to give extra scrutiny, according to this theory.

This is just the opposite, in some respects, of Green and Porter, where high price variance is characteristic of the collusion. So, here if you have markets where you can see the prices that the firms are choosing, but you believe the firms don't know very well each other's respective cost positions, then you might look for rigid pricing as characteristic of collusion. By contrast, Pn of theta has lots of price variance and is characteristic of collusion.

So, that was paper one.

Paper two, which got published earlier, with Susan Athey in Rand, 2001, addressed the issue of let's see if we can move past strong symmetric equilibria.

That's a strong assumption, so what would happen if you relaxed it.

To address that, we have taken a much simpler cost setting and say firms either have low cost or high cost. Two types. Each firm, low or high. We allow for symmetric price wars again, these wasteful transfers out of the cartel price wars, and we also now allow for asymmetric continuation values.

So, we allow now that depending on what happens in, say, period one, the firms might decide in period two to go to a pattern of future play that's more favorable to firm one than to firm two. That's something we're going to allow and you can think of that as being a transfer from firm one to firm two in the context of continuation values.

In this context, continuation values will play a couple of roles. One I just mentioned, it will play the role of a transfer, but, of course, it's not an unrestricted transfer, since you can only transfer things that correspond to equilibria of the repeated game. The second is that future play is also going to be something that you are trying to achieve efficiency in within that period as well.

Period two, for example, has sort of a double duty. As a cartel, it might be that you are assigning

play in period two to compensate some firm for admitting in period one that it had high costs. You might try to shuffle some market share toward that firm to reward it for its honesty. But also in period two itself, you want to try to be as efficient as you can as a cartel, so you have these two kinds of things that you're trying to achieve simultaneously in the second period.

Here's the result that we get: For a wide set of parameters, if firms are patient, not necessarily infinitely patient, but patient enough, optimal collusion achieves first best profit. So, they're actually able to keep price at r and assign production efficiently every single period, through a self-enforcing agreement.

So, how do they do that? Well, one way to tell the story would be to imagine that in period one, firm one draws low cost and firm two draws high cost. They get together, they meet. Firm one says, I've got low cost, really, I should be doing the producing, firm two says, okay, I've got high cost, you should go ahead, produce, charge r.

Now, why would firm two admit that it had high cost? It must be expecting some reward, some benefit from having done that. That reward, in the absence of direct side payments, would have to be that starting in

period two, somehow the equilibrium would treat firm two more favorably in compensation for having acknowledged in period one that it had high costs.

Now, if we're going to get first best profits when in period two one firm is low cost and the other firm is high cost, then we already know what has to happen. The low cost guy has to produce. So, that's the second constraint on continuation values and trying to achieve efficiency. If one of us is low and one of us is high, we know we have to assign production to the lower cost guy.

But, if both of us are low cost, or if both of us are high cost, then it doesn't matter from a joint profit perspective which of us does the producing, so we can wait around for these ties. Firm one might say, well, I'll tell you what, firm two, if you admit that you're high cost, then in period two, should we both be low cost or should we both be high cost, we'll let you do more of the producing as compensation for your acknowledgment of high cost. We'll transfer profits to you in that form.

So, the implication of this theory is that if you were looking for very sophisticated cartels, you would look for negative intertemporal correlation of firm level market shares. If one firm produced a lot in

one day, it's more likely to produce less the next day.

This prediction has some relationship with the compensation schemes used in some of the hard-core cartels that have recently been found.

Of course, this literature also suggests that there's some role for communication, although you can achieve quite a lot without communicating in the Bertrand context, and it also suggests some new perspectives on collusion and welfare. Collusion is certainly bad, as a general matter, in the presence of downward sloping demand, since firms can collude and raise price. But there is also a potential welfare benefit of collusion, which should be put into the calculations, but not over-emphasized in that the collusive firms share with society a desire to keep industry costs low.

Colluding firms are thus going to be trying to figure out schemes to assign production to the lower cost firm, which is a welfare beneficial implication of collusion, but I wouldn't want to over-emphasize that. In this model, with inelastic demand, it doesn't matter what the prices are from a welfare point of view, but in general, with downward sloping demand, the high prices that the cartel would achieve would contribute to a dead weight loss to society that would be balanced against

the potential gains and cost efficiency that a sophisticated cartel could achieve.

So, the third paper, testing your morning intellectual energy as I scroll through all of these papers, the third paper with Susan in Econometrica 2008, relaxes a different assumption of the previous papers. It relaxes the IID assumption. This assumption is very tractable, it's very helpful, but in a lot of private cost-type situations we could imagine in the real world, firms would probably have somewhat persistent private information. If a firm had low cost in period one, it's more likely to have low cost in period two as well.

So, we can't really regard the future as a completely new draw for a lot of applications that we might have in mind.

This makes things a lot more complicated. Now we have beliefs cruising around, everybody is watching each other's play saying, I saw you do this yesterday, that means probably you're this type, which means probably tomorrow you will be that type, too, and so it becomes more complicated.

As an example of the complications, consider noncooperative play. What's your analogue to Nash reversion, where do you revert if somebody cheats? It's not going to be the static Nash, that Pn of theta

function, it's not going to be that anymore, because there's going to be some dynamic signaling in the noncooperative reversion.

So, it's a more complicated environment, but it's still possible based on what I told you about the first two papers to kind of get a sense of what the main results are. So, here we will allow again for symmetric and asymmetric equilibria, and here are our results.

In this setting, if you take the extreme case, we have perfectly persistent costs, so I draw in period one a cost and I have it for the rest of time. That's kind of a benchmark case, where we have the opposite extreme from IID.

If a distribution function is log concave, so that comes back again if demand is sufficiently inelastic, and if firms are sufficiently but not necessarily infinitely patient, then optimal collusion is, again, described by rigid pricing and stable market share. So, we are back to the prediction we had in the first paper where we used the strongly symmetric equilibria concept.

So, why would that be the case, because here we do allow for the asymmetric continuation values, why would we not continue to use future favors to beat the rigid pricing scheme like we did in paper number two?

The reason is, if I'm low in period one and you're high in period one, and in the previous paper we thought, well, let's wait until tomorrow, maybe our costs would be the same and we redistribute market share then at low cost to industry profits. Well, if there's perfect persistent to cost types, tomorrow I will still be low cost and you will still be high cost, so we will never have an opportunity to redistribute market share in a way that keeps industry profit neutral.

So, while we can use these future continuation values, they don't turn out to be as helpful as they were in the IID world, because differences in cost persist through time. So, you have to wait a long time, infinitely long if costs are perfectly persistent for an evening up opportunity where you both have the same cost types.

If we step away from perfect persistence and allow imperfect persistence, we get some future market share favoring activity back into play. Fi one firm has low costs today and the other firm has high costs today, then they know that eventually they will both be low or eventually they will both be high and that will be a time where they can even up and reshuffle market shares and reward one firm for having admitted earlier on that it had high cost. So we get back some degree of this

future market share favor behavior, this negative intertemporal correlation of firm level market shares on the equilibrium path, once we step away from perfect persistence and allow for imperfect persistence.

And indeed, if the firms are patient enough, and they have to be more patient the more persistent are the cost shocks, they can again achieve a first best solution.

Let me now wrap up and then tell you about a few new directions, very briefly. This line of theoretical work just tries to say, if firms have private information, which we believe is descriptive of a lot of environments, and they don't have side payments, so they have to exchange favors through how they play the game, through their market shares, then, in a wide range of settings, where cost types are persistent, or maybe firms are not sophisticated enough to implement asymmetric continuation values, we should expect to see stable market shares and rigid pricing. For some markets, such behavior thus might be indicative of something to look at a little more carefully.

On the other hand, if firms are very sophisticated, and if their cost draws are not perfectly persistent, then they are going to use a lot of asymmetric market share allocations over time and so you

should expect to see negative intertemporal correlation at the firm level of market shares as descriptive of a very sophisticated cartel, potentially.

New directions: I'm working with a student at Stanford, Juuso Toikka, who is on the market this year, a very good student, if you're hiring, you should know about him. He and I are talking about the communication in the context of Cournot.

One other result that Susan and I had in our Econometrica paper is that if you relax the log concavity condition, then sometimes a better equilibrium for firms when types are perfectly persistent than the rigid pricing equilibrium is to have an initial price war. A firm comes in, says, I've got low cost, I'm going to cut price and I want to show you that I have low cost. Why am I going to do that? Because then in period two once they can get together and negotiate and say, you're the low cost guy, you get more market share.

So, this kind of gives rise to a perspective on price wars as preludes to bargaining. A firm may demonstrate its strength and toughness so that it can bargain better in the future over how to allocate market shares in the cartel.

Under some conditions, we get results like that. I would like to explore that more. DeRoos has a paper

in IJIO where he talks about this sort of signaling behavior offering an interpretation of behavior in the Lysine case.

And then more generally I am interested in the interaction of collusion and reputation. Once you have persistent types, models of persistent types, you can start to talk about how the reputation considerations affect cartel behavior. Cartel members may develop the reputations within the group. How does that affect their conduct? I think that's a further merging of literatures that would offer useful insights.

Thank you.

(Applause.)

DR. BAGWELL: I'm happy to take questions. Sure. Yes, sir?

UNIDENTIFIED MALE SPEAKER: Kyle, in these markets with the types, there's a particular relationship between trying to correct your reputation and what you want to try to convey to other people in the market. I know you've thought a little bit about auctions, which I would like to hear a bit more about.

Can you think about it in the same way where the information that you would want to convey to the other people in the market would be the same when instead of you wanting to cooperate with them, you actually want to

compete with them, so the reputation has sort of a different feel for it? I'm not sure if you can answer that.

DR. BAGWELL: Well, I mean, you can certainly recast everything that I've described as a Bertrand kind of model with public prices, and some might think of it as some product market, but you could think of it alternatively as a procurement auction wherein bidders are repeatedly bidding for objects and they have their valuations instead of their cost types, and they're going to try to collude in this context.

And then everything would kind of map over to a story where you would either have valuations that were independent through time, or persistent through time, have the opportunity, perhaps, to talk ahead of time about how you're going to bid, and the framework could be applied there.

And then on the general notion of reputation, so far, we've had a few results, like pooling is optimal in some cases and you can get first best collusion with imperfect persistence if you're patient enough, but I think this price war idea that we addressed some in the Econometrica paper that would be nice to emphasize more and build on more in future work. It's really suggesting you can do something demonstrative to say, I

am a tough member of the cartel and I need to be treated as such in the early phase. That's a story that seems realistic, but it also clearly requires some persistence in the type space for it to get traction.

UNIDENTIFIED MALE SPEAKER: Kyle, we have a question, you emphasized as one possible screen for collusion, variance in price. Have you also looked maybe by the exercises into price level? I know it's not really very much use, but from data I've seen on a particular intermediary market, it's something that often seems that there's some price floor where maybe this is for other reasons, but you can also think of it as a way to collude on the price.

My feeling is that maybe your theory would also support something like what is the appropriate data, something like a skew measure, which would come up as a potential measure of screening, I think to compare it then with information about cost using the authority we have.

DR. BAGWELL: I think that's a good point. So, you're saying that sometimes you might see a cartel where there's seemingly a price floor and the agreement is if you want to price higher, go ahead, but don't go below this floor, as that would be interpreted as a deviation.

Yeah, I think that's a very good point, and I think there are models I can imagine that would emphasize downward sloping demand as opposed to the inelastic structure here where that might well come out as optimal. Susan has a paper on monetary policy with Atkinson and Kehoe which has a ceiling, kind of a similar structure, and there's some work in trade policy that I have with Bob Staiger that has a similar structure.

So, I would imagine that that kind of result could arise under some conditions in a collusion model with downward sloping demand and that it would be a good thing to analyze further, particularly if that's what is actually happening in practice, or seems to be.

UNIDENTIFIED MALE SPEAKER: I'm always nervous about screens that have to do with the variance that you see in price, because obviously in a competitive market, you see no variance, right? And I'm also nervous about who the audience is for that kind of advice, because the audience is not always capable of sorting through, well, under this condition, with this kind of a model, the screen means this, but with a different kind of a model, the screen might mean something else.

Do you have any kind of advice or sort of qualification that you would put on using those kinds of

screens to detect collusion that would be advice that would help policy?

DR. BAGWELL: Well, first of all, I would like to thank you for giving me the opportunity to emphasize, again, that these screening tests that I am describing should be viewed with great caution, right? They're neither sufficient nor necessary for finding collusion, but just perhaps informative in certain settings.

So, the sort of settings would be ones where firms costs are bouncing around a lot. Then the price variance would be high in the competitive case, which was the Nash case that I just described, and low in the collusive case.

If firms costs are very stable through time, then as you actually said, a competitive equilibrium would have very low price variance, and you wouldn't want to get too excited about that and pursue antitrust enforcement there.

There is an interesting IJIO paper that I'll just mention by Luke Froeb and others, I can't remember all of the co-authors, that considers a variance screening test for collusion. But I really don't want to over-emphasize screening in this talk other than to just say that low price variance is something that this theory points to as characteristic of optimal collusion

in some settings. It could lead to a cautious assessment that maybe it would be worthwhile to look at a certain market a little more carefully. But as a general matter, low price variance is not a sufficient condition for concluding that there's collusion, and it is also not a necessary condition for concluding that there's collusion.

So, I do not want to emphasize that more heavily than it deserves.

UNIDENTIFIED MALE SPEAKER: Would it be interesting to add, I'm just going to try to project, would it be interesting to add to these models something to the effect of cost is a function of an appropriatable technology that then firms could agree to share?

DR. BAGWELL: I'm sorry, they have a technology that they might share?

UNIDENTIFIED MALE SPEAKER: So, perhaps I've developed some innovation that's currently privatized and until somebody leap-frogs me, I'm going to be the lowest cost firm in the market, but to preserve the cartel, or need to maximize profits for the cartel, it would be efficient for me to share it.

DR. BAGWELL: So, yes, that would be very interesting, and we haven't introduced some novel incentive issues, we haven't looked at that, but that

would be definitely the kind of question that one might use this framework to think more about, for sure.

DR. ROTHSTEIN: Thank you very much.

DR. BAGWELL: Thank you.

(Applause.)

DR. ROTHSTEIN: So, we're setting up for our next paper session, Asymmetric Information and Consumer Choice. Just a couple of quick announcements. One, I would like to thank Marianne Bertrand for helping to set up this session, but she can't be here today.

So, Erez Yoeli from the Division of Consumer Protection is going to chair. There is some symmetry there, Marianne is at Booth Chicago, and Erez just got his degree from Booth Chicago. His diploma is probably still wet. Where is Erez?

UNIDENTIFIED MALE SPEAKER: He's coming right back. I'll go grab him.

DR. ROTHSTEIN: I forget, he's been here maybe two months. Four months, tops. There's also one substitution, Devin Pope is here, instead of Justin Sydnor. Justin couldn't make it, but we have his co-author here. So, I am going to turn it over to Erez.

DR. YOELI: Hi, I'm taking Marianne's place, and I'm going to, without further ado, introduce Liad Wagman, he is assistant professor of economics at IIT,

the Illinois Institute of Economics in Chicago, and recently finished his Ph.D. at Duke. He also has a computer science background, which is a little rare.

Liad?

DR. WAGMAN: Thanks, Erez.

I am going to talk about online privacy today, and online privacy is a loaded issue. There are many facets to it: Who collects the data, what data is being collected, how is it being used, who has control of the data, and I don't want to pretend that we're able to tackle all of it in this short paper.

We basically take a slice, and the slice we take is data that's being collected after transactions. It's basically how much consumers were paying for a certain good, and we're giving consumers some form of control over this data. Perhaps it's best if I start with a few examples.

This is a coupon for buy.com that gives you $10 off a purchase of $200 or more. The terms and conditions state that it's valid for first-time customers only. Only you can become a first-time customer by just getting another email address.

It may not seem very attractive in this case, because the purchase has to be at least $200, so it's a five percent discount, almost nothing. But a little

while ago, Google Checkout came out with a promotion that for a short time, when you sign up, you get $10 off a purchase of $10 or more. This promotion happened to be applicable to buy.com, a site that doesn't charge tax in most states, has free shipping on most items, and has a lot of $10 offerings. Think of flash drives, things of that sort.

So, you could get a $10 purchase, essentially for free, if you signed up for Google Checkout. Now, obviously you would have incentive to do that many times if you wanted to stockpile some flash dives, maybe sell them on eBay, make some money.

So, initially if you just change, get another email address, sign up again, get the discount again. Well, Google quickly caught up to that, although they have an extensive budget for promotion, they didn't want a few consumers to hog the entire budget. So, they limited it initially by email address.

Well, that didn't work out too well. You could just create another email address. So, then they limited it by address, home address. That didn't work out too well either, because you could put slight modifications, perhaps add a unit number to your place, even though it has no unit, maybe use your work address, so forth.

So, then they limited it by credit cards. If you used a credit card, you could use it again to get the promotion. Well, you could use another credit card or you could even generate a virtual credit card. It's a one-time use credit card derived from your real credit card, which you can generate in seconds. Just go on Citibank Online, if you have a Citibank credit card, you go and generate a virtual credit card, it's for one-time use. Feel free to copy the number, you can't use it.

So, that didn't work out either. So, Google, being pretty innovative, started tracking people by IP address as well. So, they had some way to account for public libraries and so forth, but if you used your IP address at home and based on your IP address, they could tell that you were a private user, they wouldn't let you use it again.

So, what we see here is on the one hand you have the firm trying to track these first-time consumers, on the other hand you have consumers trying to be anonymous and taking advantage of these services, multiple times, and it's quite interesting how that interaction goes.

Perhaps more related, I went a few months on dell.com, and I clicked on medium/small businesses, and I went to the printer section. These few printers showed up on the front page. Do you see that they have

certain prices, certain models? And then I erased my cookies, closed my browser, opened it again, went back to dell.com. Then I clicked on large businesses, instead of medium or small businesses, and went to the printer section again.

Surprisingly, I got the same printers, order was shuffled a little bit, prices were at least 20 percent higher.

A more famous example is the Amazon DVD price experimentation that took place in 2000 where they charged people different price based on their past searches, or just on their viewing behavior on the site. If you looked for a StarWars DVD, you would get charged more for a StarTrek DVD. If you bought a StarWars DVD, you would get charged more for a StarTrek DVD, up to 40 percent more.

So, price discrimination is pretty common in electronic commerce. Consumers, at the same time, are not helpless. They can work to become anonymous, to maintain their privacy, to maintain their anonymity, and they can circumvent it.

Currently sellers' practices as far as price discrimination goes mostly follow voluntary guidelines. There are no strict rules or policy. But, we are moving towards some policy, it seems.

About two years ago, the Commission came out with guidelines on how sellers should set up their sites, and those guidelines mainly emphasized transparency in disclosures, making disclosures easy to understand, and consumer control, giving consumers the opportunity to choose to opt out of having their information collected. I'll talk about opt-in a little bit later, if I have time, but it's almost the same picture.

There's other movement towards privacy, towards litigation here. There's the Online Privacy Bill of Rights proposed by Edward Markey in Congress. There's similar litigation that already exists in telecommunication, there's the CAN-SPAM Act of 2003, and there's more. The general direction here is to make it easier for consumers to maintain anonymity.

So, just to summarize the key differences we have here from traditional markets, it is already easy for consumers to become anonymous. At the same time, sellers can easily store and use consumers' date, should the consumer not be anonymous.

The overall question we're trying to tackle is is it easy and desireable to obtain anonymity and is it clear who benefits and loses from that?

To do that, we study a very simple game, or at

least we thought it was simple when we started. There's one firm, potentially more firms and many consumers, and there are three main parts. The first part is some past purchases disclose information about consumers' valuations. In the second stage, consumers make anonymity decisions, they decide whether to work to become anonymous again.

In the third stage, if the firm has some information about consumers, it can use this information to price discriminate. The firm has some identified consumers that purchased and didn't become anonymous, and it can charge them a little bit more because it knows that they were willing to pay a certain price. There are also consumers who are anonymous, who either didn't purchase in the past or maybe purchased and worked to become anonymous again.

So, when you work on a theoretical model, the main goal is to have a simple model, maybe have some closed form solutions and maybe get some clear-cut predictions. Well, the model turned out not to be so simple, the solutions were not closed form and the predictions were not clear-cut. That turned out to be actually good, because we are able to say that even this simple framework, with black and white economic analysis, it is not clear whether more privacy is good

or bad. I'll show that more precisely in a little bit.

I am going to follow Sofia's action from yesterday and sort of assassinate the literature review slide, but I just want to emphasize that it's a very interdisciplinary topic, it goes form psychology to computer science to business to marketing to economics, it's all over. The literature is pretty extensive, and although it's a fairly recent topic, it dates back to the seventies, with the announcement of intertemporal price discrimination, and much of the intuition has gone from there, so that still holds.

So, the model we have is fairly standard in this framework. There were two purchasing periods. The firm produces a nondurable good, with no marginal cost. There is a continuum of strategic consumers with mass one, and the timeline is as follows: Consumers first privately realize valuation, the firm sets the first period price, consumers make their first period purchasing decisions, those who purchased need to decide whether to work to become anonymous again or not. Then we have the firm setting prices, and then consumers make their purchasing decisions again.

So, and I'll show an extensive form view of this in a second, because I understand I went quickly over the model. In each period a consumer has unit demand,

valuation is drawn from continuous distribution on the unit interval. The valuation is private information and we assume it to be the same in both periods. The cost of opting out or maintaining your anonymity is assumed to be c and is assumed to be the same for consumers. The results were somewhat robust to making this different for different consumers, and perhaps I can show it a little bit in a few slides.

This is the extensive form view of the game, or a sketch of it. The firm decides its first period price, consumers' valuations are realized. I flipped the order here, but it doesn't matter, it just makes it easier to draw. Then consumers make their purchasing decisions. They can decide whether to buy or not buy. Those who buy decide whether to opt out or not opt out, and then the firm has two information sets. It has the information set for identified consumers, those who bought and didn't opt out, so it knows they were willing to pay a certain price for the good in the first period. Then it has the information set for unidentified consumers, or anonymous consumers. Those are consumers who either bought and opted out, or didn't buy at all. Then it sets prices, two prices, one to each of these information sets. Then consumers, again, make their purchasing decisions in the second purchasing period.

Quick overview of the results. Given that the firm cannot commit to future prices, the assumption that could be reasonable in some settings. So with the printer example at Dell, it could go to the extent where Dell supplies you with a driver, that you could install on your computer when you set up the printer. This driver can tell Dell your usage patterns, and when your printer cartridge is empty, when you need to order some toner, it can supply you with a link, a personalized link to Dell where it can tell you a price, based on your usage patterns, based on the price you paid for the printer when you were buying initially. So, the assumption here is the firm cannot commit to future prices when you make your initial purchase.

Given some technical assumptions, the firm's profit is nonmonotonic in the cost of opting out, the cost of being anonymous. Profit is highest when this cost is zero. That is to say it is highest when consumers have full control over their anonymity.

Consumer surplus, on the other hand, may increase their cost of maintaining anonymity, but only up to a point. At that point, it goes down. So, the ability to opt out at some costs is actually good, but if it's too easy to opt out, it may not be good.

There are similar results with social surplus,

almost the same as consumer surplus, at least to some extent. Some preliminary observations and results here are that the socially optimal outcome is for all consumers to purchase in each period, because the marginal cost is zero. That's just how the model is set up, that would be the first best. But that's never going to happen in the model with one firm, because the firm will set price greater than zero.

Now, if there's no consumer recognition, the firm just sets the monopoly price in each period. It's just the one-shot game played twice.

If the firm can commit to future prices, and opting out is impossible, it's too costly, the firm would like to commit to the monopoly prices.

So far these are the results from the existing literature, now we add to this. If the firm can commit to future prices, and opting out is possible, the firm would still like to commit to the monopoly prices. So, that's the full commitment outcome. The full commitment outcome is the same, as in the case where there's no consumer recognition.

So, to characterize the model, remember the consumers can opt out at the cost, c. If this cost is zero, meaning anonymity is free to obtain, I can buy and I can instantly become anonymous again. All perfect

Bayesian equilibria have the following properties: Prices are equal to monopoly prices, this is the full commitment outcome. Consumers with valuation above price purchase in both periods, all of them stay anonymous, there is no recognition outcomes. This is exactly what the firm wants.

So, if consumers can costlessly become anonymous after they purchase, the firm gets what it wants, it gets the full commitment outcome. This is sort of counter-intuitive right there, because if you think that consumers have full control, why shouldn't they? I hope I can get some intuition on that.

So, the key thing is that opting out here is associated with a negative externality on other consumers. Individually, a consumer wants to opt out to have access to cheaper prices. But doing so, anonymous consumers pay more, because the firm targets more consumers with high valuation in the anonymous consumer pool.

And these are similar situations to Prisoner's Dilemma, Tragedy of the Commons or even Braess's Paradox in the consumer science. I will talk about Braess's Paradox a little bit in a few slides.

So, let's start with the last stage of the game. We are in the second purchasing period, the firm has

some information about consumers, and now the cost of opting out is greater than zero. There are some consumers that are repeat consumers, they already bought. Those are above the cut-off of the valuation v, and a fraction of them chose to opt out, a fraction of them are still identified.

Let alpha be the portion of consumers that maintain their anonymity, that opted out, out of those that bought in the first period. So the firm targets these two information sets, unidentified consumers, and repeat consumers that are still identified, and it sets prices optimally to these two consumer briefs. This is just a mathematical question where F is the distribution, v is the cut-off point, alpha is the proportion of consumers that opted out, and it sets the price optimally to identified consumers. We have the typical ratchet effect on the literature here where the consumers who are identified pay more, because the firm knows that they signaled a higher willingness to pay, that they were willing to pay the price in the period previously purchased.

So, this is pretty standard. Now, in the second stage, consumers choose their anonymity. This is very similar to a network flow problem, and this is where Braess's Paradox comes in. Consumers will opt out, as

long as there is an incentive to do so, and as long as the price differential is greater than c, greater than the cost of opting out.

From that, we can derive the portion of consumers that choose to opt out. So, Braess's Paradox basically says, we have this first period in which consumers can buy or not buy. If they buy, they are identified; if they don't buy, they are anonymous. Then there are second period purchases.

If we add the edge from identified consumers to anonymous consumers, where consumers can opt out, there will be a flow of consumers moving from the top segment to the bottom segment, as long as it pays off to do so. But the addition of that edge actually ends up hurting everybody.

It reduces traffic at the bottom flow to the point where the advantage dissipates, where nobody benefits from going in that direction.

So, that's exactly what happens here. Consumers will opt out, as long as there's an advantage to doing so, because of the cheaper prices later on. They will continue doing it, and the advantage will dissipate. There will be no advantage to being anonymous anymore.

In the first stage, the firm sets its first period price, optimally, taking everything under

account. It's just a mathematical expression that comes out of it, and this is what ends up happening in equilibrium as a function of the cost of opting out. When this cost is zero, everybody opts out, so alpha is equal to one, there's no recognition.

As this cost increases, less and less consumers opt out. At some point, c-hat, nobody opts out anymore, but prices still change. The reason being that the firm had to decrease prices to get consumers to buy, because the consumers would know that once they bought, they disclosed some information and there's some cost associated with getting rid of that information that was collected.

So, the firm had to lower prices to get consumers to buy. It lowered them to the point where no more consumers were opting out. Now we can increase prices to better price discriminate. That's this interval between c-hat and c-upper bar. Prices change, even when no consumer opts out.

And once we reach c-upper bar, there are no more changes, the collusion is static. It's the case where opting out is too costly, nobody is going to do it.

All right. So, let's see the consumer surplus in the uniform case, as a function of this cost of opting out. As you can see, it's low risk when this

cost is zero, it increases with this cost because the firm lowers prices to get consumers to buy, up to c-hat. At that point it goes down a little bit, because the firm further price discriminates.

The firm's profit is also nonmonotonic, it goes down sharply in this cost because the firm lowers prices, but eventually, as fewer and fewer consumers opt out, it slightly goes back up and the firm is able to recoup some profit through price discrimination.

We found through these two we can get some measure of social surplus here. Social surplus initially goes down because there's a lot of waste from this effort being put into becoming anonymous, but eventually picks up as more and more consumers buy due to lower prices. Then it goes up all the way to c-hat where it drops because the firm increases prices. It further price discriminates once nobody opts out.

So, what we get here is that some privacy, specifically making it, say, between c-hat and c-bar in the uniform case, is actually good. Lowering it to c-hat is good. But lowering it too much is actually bad.

I wanted to give just a direction here for what happened if the firm could commit not to charge identified consumers more. So, that's Amazon, after the

DVD pricing experiment. It basically committed not to charge more, but it still discriminates downwards. It offers coupons, incentives to identified consumers to stay identified.

So, what would happen there is that we would have some gain where consumers' valuation could fall on market process, the process could be known, but the current and past valuations are not exactly known, but private information, and the firm learns about these valuations through consumer purchases. It's something similar to a loyalty program. Prices have to be low enough to incentivize consumers to buy using their membership account, without manipulating the program.

So, that's what would happen if the firm could commit not to charge identified consumers more, which is what is happening in the case of Amazon.

We could also extend this setting to a case of competition, where firms are competing so that in this model I have two firms. The market leader is selling a brand name product and a follower is selling a generic product. The brand name may be valued a little more, according to some parameter gamma, and there could be three regimes where there's no recognition, where there's asymmetric recognition, perhaps only the market leader can recognize past transactions and past

consumers, and there could be full recognition where both of them could recognize consumers.

The result here is that it's still ambiguous whether social surplus is higher or lower under a fuller commission or lower commission. Consumers, though, are better off under recognition than no recognition.

So, just to summarize, we found that in the base model, we had maximum profit for this firm when the cost of becoming anonymous is zero, because the firm is able to obtain the full commitment outcome. Making privacy or anonymity easy to maintain can increase, but also decrease welfare and consumer participation. So, it's good to some extent, but doing too much of it is not good. It basically comes to show that even in this simple framework, it's not clear-cut whether more privacy is good or bad.

We have nonmonotonicity in social surplus, in profit, and the various extensions here that we could consider to commit to competition.

DR. YOELI: Next we have Alessandro Pavan from Northwestern's Economics Department.

DR. PAVAN: So, sorry for my sore throat, it's terrible.

So, let me just say up front that this is an interesting paper, so I very much enjoyed reading it.

It's thought provoking and I think it touches on an important topic.

Let me give you a very, very quick overview of what the paper is trying to say and then offer some suggestion and comments regarding future work.

As Liad has indicated at the beginning of the presentation, privacy is an extremely important aspect of online transactions, not only of online, but of any transaction. The effect of privacy on consumer surplus, reduced surplus and welfare, quite surprisingly, are very much largely not understood, or not properly understood. That explains why, I think, there is a significant need for more theoretical work and this paper can contribute nicely in this direction.

So, the model is extremely simple, but nevertheless it offers some very insightful points that I think should help reach a better understanding of what are the effects of privacy in a variety of settings.

So, the model, I know you have seen how it works, is essentially couched in a situation where you have a monopoly selling identical goods over two periods. You have a continuum of buyers with unit demands in each period, and an important assumption is that buyers have additively separable valuations, and so the terms of trade in respect to the second transaction

do not depend on the terms of trade in the first transaction, but we are affected by the first transaction by the information which is disclosed to the seller.

Buyers can choose at some cost to remain anonymous, and this cost is sort of proxy for the cost of, say, deleting cookies or engaging in an effort that prevents seller from recognizing prior purchases.

Consumers have private information in the valuation, as well as the past transactions, and monopolists cannot commit to future prices, that's an important assumption.

So, let me first say that it's sort of a common misperception around privacy, so first is due to the fact that privacy means different things to different audiences, but also, if you define privacy, in this sort of narrow way, even in that case, I think there's some perception, or a misperception, I should say, out there, that consumers should benefit from privacy. The argument being that we know that privacy, or private information, I should say, is what entitled the consumer to enjoy informational rights. So, if privacy is gone, then the seller will exploit it, and appropriate more surplus.

So, (b) the seller on the other hand is

benefitting from the consumer's inability to protect privacy, because it can appropriate more surplus. If you combine (a) and (b), you get the policy prescription in favor of privacy-protecting laws with less financing and generally no governments or regulators assure that the equal measures aimed at reducing cost of privacy.

So, this is essentially what often you read in the public press. On the other hand, a large body of research in industrial organization, the Contract Theory and Information Economics that dates back to Coase, Baron and Besanko, Vickers, Laffont and Tirole, just to mention a few, have indicated that actually firms would love to commit and not price discriminate.

So, this seems to indicate that maybe the benefits and who suffers from privacy is not as clear-cut as often indicated in the popular press, that's why we need more theoretical analysis.

So, let me also say that you can approach privacy, essentially from two different perspectives, so it better reflects what Richard has been doing over the years, and the two perspectives actually reflect who controls privacy. So, the early literature assumes that privacy policies are designed by the firms. So, Coase's early work, effort in this area, I had some work with Calzolari, and I also had some paper work in this area.

So, the early literature assumes that privacy is controlled by the firms, and then this sort of benchmark that we established in the work with Calzolari essentially says the following: Suppose you have streams of sellers who contract with the same buyers, and suppose the downstream sellers are not personally interested in downstream trade. So, if you don't engage in any sort of profit sharing, and suppose we have the consumers who value upstream products the most compared to other consumers also value downstream products the most, which technically corresponds to assuming that consumers' valuation had a constant sign of the single crossing condition, and finally assumes that the consumers have additively separable preferences, exactly as in this model, and then you have a sort of benchmark, a theorem that says that upstream sellers commit to full privacy. That is so even if they could not make money by selling information to downstream sellers.

So, if privacy is controlled by the firms, and if these assumptions are met, then the optimal privacy policy is full privacy. That is not so obvious, because there are essentially three effects of the information sharing: One is if you disclose information say to a business partner, but you can make money out of it, another is if you can sell the information, and another

is a rent-shifting effect, so if you disclose the information, then you can reduce the incentive for information in a downstream relationship, leading to a higher level of trade and even permitting your customer to take advantage of that.

So, on the other hand, privacy has an effect on incentives for information-revelation upstream, and what the theory essentially is saying here is that when the conditions of the theorem are met, the first two effects, even if positive, are more than offset by the implication that privacy creates on incentives for information-revelation upstream to the point that from the perspective of the upstream firm, the best is to commit to full privacy.

So, the paper presented today departs from this earlier literature by assuming that privacy is controlled by the consumer, at a cost. When the cost is zero, then the key result is that there is a coordination failure, where consumers in equilibrium opted for full privacy and that's exactly what firms want. That permits essentially the firms to use the same monopoly prices in both periods, which is consistent with Baron and Besanko's prediction; that's exactly what a monopoly is going to commit it would do.

So, this is essentially a form of Prisoner's

Dilemma, or coordination failure, as was mentioned a moment ago, where consumers don't even realize the effect of the privacy choices on equilibrium prices.

So, because we don't do it, then that would be the worst possible situation, where customer surplus is minimum. So, starting from this point, then the rest of the paper examines the effect of a variation in the cost of privacy on consumer surplus, producer surplus and total welfare, which has been nicely highlighted by the argument, so I will not repeat here what is going on.

So, I just can see from the picture, consumer surplus is not monotone in the cost of privacy, but is minimal when the cost of privacy is zero, because that's exactly the situation where firms can engage in essentially monopoly pricing in both periods.

The other side of the coin is where the profits are decreasing and the cost of guaranteeing your privacy, and social surplus is also more monotone now, as you can see.

So, let me conclude with a couple of remarks on privacy and a few questions, or suggestions for future work.

So, maybe the advocates of privacy in the popular press are afraid that the consumers are not as smart as in our models. If that is the case, then there

is a good argument for favoring privacy by reducing the cost of enforcing it.

There's a number of reasons for why on the other hand privacy may not be desirable, which is missing in the model, is the complementarity and substantiality between goods or more generally, if you think about privacy as undermining the possibility of a downstream seller to sort of tailor the terms of trade to the needs of the buyer. You have to be careful when you advocate for full privacy because essentially you are neglecting that information sharing from upstream seller to downstream seller could also create welfare given buyer needs.

As far as suggestions for future work, one of the first things that came to my mind when I was reading the paper is that if we really wanted to understand whether privacy is desirable or not, and whether regulators and policy makers should try to reduce the cost of privacy, I think actually we need a model where competition actually is formalized in a little bit more exhaustive way than in the paper. I think it's really important to extend the work in order to capture competition beyond what has already been indicated. So, that's I think a first consideration for future work.

I would also like to see an alternative to the

pooling equilibrium; personally I think there should be strategy. I believe that should be forbidden. So, maybe there is a way of getting essentially the same without requiring this mixing on our consumer side, so you model privacy as a continuous choice as opposed to a discrete one.

I would also like to suggest that I think it would be beneficial for the reader to illustrate what happens off-equilibrium. So, I think we are accustomed to analyzing the equilibrium, but to shed light on what is really happening, you really need to consider what happens off-equilibrium.

So, you start from say, for example, that the prices are exactly the same as under the case where cost is equal to zero, and assuming now if cost goes up, and then show exactly why the same prices cannot be sustained anymore. So, we start off with the economics of the new equilibrium, where often that is useful in shedding light on what is happening.

Lastly, I would like to see a better explanation for the costs of privacy. So, you mentioned some of them, but I was not convinced of what you really had in mind when you talked about the cost of enforcing it. Some empirical work is always welcomed.

DR. YOELI: Thank you. Next on the list is

Marco Ottaviani, he is a professor of management and strategy at Kellogg.

DR. INDERST: For those who know Marco, that's not me, actually. I am his co-author and let me first find my slide before I say two or three words in introduction. I think I missed sending in my CV, et cetera. I was reminded that I didn't complete that.

One or two words about me, I am professor of economics and finance in particular in Frankfurt, I used to be a professor of economics at the London School of Economics. This paper, joint with Marco, is part of an overall joint research effort to work on economics of advice, in particular economics of advice in retail finance, and then asking a shift of the regulation.

As a part of motivation for this particular paper, let me read from the roadmap of the U.S. Treasury to rebuild financial supervision, presumably from scratch. It says that impartial advice represents one of the most important financial services for consumers. Mortgage brokers, as an example, often advertise their trustworthiness as advisors on difficult mortgage decisions. When these intermediaries accept side payments, they can compromise their ability to be impartial. Consumers may retain the faith that the intermediary is working for them and placing their

interests above his or her own, even if the conflict of interest is disclosed. Accordingly, in some cases, consumers may reasonably be mistaken when they rely on advice.

Now, this quote summarizes a bit the key ingredients that are going to be now modeled, which is the role of financial advice going to be key. Second, financial advice can be compromised, potentially, by side payments, commissions. Thirdly, consumers, customers, may not adequately anticipate and take this into account in their decisions.

Let me add a little bit of extra details on each of these ingredients, before I then come to the model. The first being the role of financial advice. My starting point here is the theory of household finance. Household finance market economic theory, in particular, has tried to understand customers' investment decisions from the perspective of an active decision-maker who requires information and then potentially, due to behavioral advisors, makes some decisions which are wrong, but other decisions which are better explained in the standard paradigm, et cetera.

Now, this is quite okay for some individuals, I would say, but it misses a key point or a key aspect of the industry. That is that many of the financial

decisions are basically not initiated by the customer, but initiated by the seller. Financial products are often sold and not bought, and they are bought after getting substantial financial advice.

Our paper has a lot of evidence of this in our footnote and introduction, which I would like to mention one or two of these. There was in Europe a big survey in 2003 asking a lot of questions about customers' attitudes on different products, et cetera, and also asking customers across different European countries whether they expect to receive financial advice from institutions, and 90 percent of customers in Austria, Germany, et cetera, do expect to receive financial advice and to overwhelmingly trust financial advice.

In the U.S., there is also evidence that the financial advice is very important, for instance in the mutual funds industry.

Now, how do customers pay for financial advice? Even though there are different forms of payment out there, the most common form still seems to be indirectly, this is through commissions or distribution fees that the advisor gets paid, and ultimately, customers have to pay for them through being charged higher fees, higher loads.

The same is happening with mortgages and yield

spreads. Of course, that may compromise the advice of the financial advisor and this is also recognized in the industry. For instance, this is a survey of European professionals that 64 percent say, well, we believe our fee structure drives our sales to customers and not suitability to customers.

And the FSA, for instance, has drawn a lesson from this, and basically, although this is still in process, it tries or plans to stamp out commissions completely from independent financial advice by saying, okay, you have advised the firm, then it needs to be paid directly by the charges and it can no longer accept commissions. That's still basically proposed reform and, the industry, of course, reacted quite aggressively saying that, well, if that's going to be passed as regulation, 30 percent of independent financial advisors are going to be pushed out of the market; I don't know whether it's bad or good.

For U.S., generally there is very limited, very limited evidence on how financial advice is compromised out there. For the U.S. mutual funds industry, there is indirect evidence suggesting that funds sold through brokers underperformed those funds which had people collect themselves and that higher fees basically drive distribution.

Let me give you, before I get into the paper, additional direct evidence that we currently have in our paper, together with joint authors from Frankfurt, a paper called "The Dark Side of Financial Advice." Go to the home page and it also gives a little background of why in Germany advice is given mainly by large banks.

Now, what's unique in this data is that we have a portfolio of information from customers of the bank and a detailed customer survey, so we know who relies on financial advice and why, et cetera. In addition, we know what the bank makes per customer on security transactions, and also what fraction of the products people are buying have been incentivized. Incentivized means basically that you have to sell or that's what you should sell in this particular week or month.

What we found is that those who learned from advice were often less informed, less educated, they are older and, which is very significant, they are not self employed. Those who rely less on advice question the advisor's suggestion, which is a knowledge issue, and would be not relying solely on the advice and would be willing to question the advisor because they are more aware of the conflict of interest.

There are a lot of questions there, and data which support that people who do and do not rely on

advice have a different conception and perception of the conflict of interest.

Last line on this evidence, well, what does reliance on advice mean in the end? Well, it means in the end that if you rely strongly on advice compared to not so strongly, then you end up with a much larger fraction of incentivized products in your portfolio. And more strikingly, customers who say that they rely strongly on advice end up generating for the bank 20 percent more revenues over two years, and that's because of higher trade.

A lot of trading is basically instigated by the advisor. It's not due to over-confidence of the customer, but we can really document that it's typically the advisor who initiates a talk which leads them to a potential turn of the portfolio.

And, of course, here with this large bank, how is advice paid for? It's not paid for on the basis of each sales talk you have to come up with a fee that you have to pay, but maybe you pay indirectly through the loads and therefore the higher revenues.

Now, what rationalizes this form of paying for financial advice? It could be that it's because of legal duties, different fiduciary duty, and then we have something to say on this also in the paper, but I cannot

explain it for the next 15 minutes.

Instead, I would like to highlight the two different reasons that we identify for indirect payments for financial advice. One has to do much with strategically naive customers. Basically these customers do not see and anticipate fully the conflict of interest between them and the advisor. That seems to be also happening in all the papers on the last line of the slide. There is evidence from experiments by Loewenstein, for instance, we are running experiments in Germany on this, where we can really concede that some customers in persuasion games seem to be more naive than others, but there is much more work needed here. But we use this as a key ingredient to first power our model, which is some customers are naive. We find on equilibrium, these customers who are naive are not charged any direct payment for advice. The reason for why you let them pay indirectly is basically an exploitive mechanism. It allows you to exploit these customers as much as possible to extract profits from them as much as possible.

A cap or a ban on indirect payments and commissions would then be consumer surplus enhancing.

However, we show that with wary customers this can backfire, because with wary customers this may be

the second best vision outcome. That comes a little bit later, and I'm going to talk in a little bit about other policies, other policy possibilities, other than regulating commissions, for instance, like health warnings and disclosure.

Let me now for the next ten minutes, last ten minutes, go into the model. The baseline model is very similar. We have a customer who has to choose between two options, A and B. You can call A an advanced, a premium product, for example an equity investment, on which the bank specifically earns more, and B a basic product. B is basically the informed option of not buying.

Now, customers come in two types, which we also know as A or B. The utility of the customer is not derived from two options and depends on his type and the option, and if the option basically matches your type, we assume, a high return vh, and if the options doesn't match your type, we derive a low return vl. It's basically about finding the right fit for you to fit the profile.

Think about different investment opportunities which have different risk return profiles, different tax possibilities, and different customers with different income status, different risk aversion, et cetera, et

cetera. So, it's about matching customers to particular products. The prior belief that A is a better match is denoted by q0, and that's all the customer knows.

This is where the advisor comes in, because the advisor can generate better information, he knows the products better, and by exerting some effort, he can also become better acquainted with the customer situation. This allows him then to generate additional privately observed information, which we then summarize by saying, okay, he will have then have a better informed posterior belief q, which is the probability that for this customer, which just entered the door, product A is more suitable.

Now, we model here the informativeness of the advisor's posterior belief directly by placing an ordering on the distribution of the posterior release, because it's a very convenient way. We say that more effort will induce a mean preserving rotation of the standard distribution of the posterior beliefs.

So, think about the extreme cases where the advisor doesn't observe anything, in addition, and basically his posterior belief will focus all of the mass in q=q0, which is the prior belief. So, his posterior beliefs will be the same as the example.

On the other hand, if the advisor puts in a lot

of effort and gets precise information, and from exact perspective, his posterior belief will be either q=0 or q=1. So, he will know that either B or A, which provides the best fit. The mean preserving rotation just makes this shift more monotonically.

Let me now after having laid out the technicalities and technology of the paper, come to the contracting end of the game. Now, in t=1, the contracting take place. The product provider, the seller, offers an advisor, the intermediary, a fixed part, Ta, and the commission, which is then paid, if later on a customer purchases product A.

And at the same time, the product provider specifies a price for the product. It's a very general framework, but I have motivated this with respect to the finance.

In option B, in this paper, with other papers we do a different connection, the option B in this paper is basically not provided strategically. We focus really on option A. So, B could be option of not purchasing, or purchasing a competitor-provided product, and that is why we specify that the commission that the intermediary collects in B just covers his common handling costs, k. Tb=k. The price that the customer pays just cover the handling cost plus the production cost. So, we are

really focusing here on the strategically provided product, A, the advanced product.

Well, this is what's happening contracting-wise in the first period, and now the game unfolds. The advisor can now himself set a fee for advice, f, and then the customer comes in and can accept his contract, pay the fee. When he pays the fee, the advisor gets paid for his efforts, observes the additional information, then gets the posterior q, which is the probability that A would fit this customer, and makes the recommendation.

So, this is at this stage a cheap talk game and we are focused here if it exists on the informative equilibrium.

Finally, in t=5, the customer decides whether to take the advice or not, and pays a realized.

Now, let's look at the advice stage. The advisor's position, driven from posterior q, whether to recommend product A or B, and the assumption that the customer follows his advice. The advisor must take into account costs "rho", strictly positive costs, that depend on the probability A is in fact the wrong choice.

So, these costs "rho", which the advisor then has to appear when an unsuitable match basically happens in the end, then you could think of as reputational

costs, liability costs, et cetera, et cetera. We are very agnostic about this in this paper, because "rho" does not for this paper constitute, say, a policy variable. This is what "rho" could come from different sources.

Now, given "rho," these costs that the advisor has to pay, when the match was not perfect, and given the commissions, what is now the advisor's decision what to recommend? Well, the advisor now recommends A, if the following in the quality holds, which is if the commission Ta minus the handling costs, minus the expected costs, transportation costs, et cetera, when A is not the right fit, which is (1-q), multiplied by "rho", when this is at least as large as the symmetric term for product B.

And if then q gives rise to a cutoff q*, such that if the advisor is more optimistic than q*, so if q is larger than q*, he will be more optimistic about A, then he will recommend A, otherwise he will recommend B.

This is what the advisor does at this stage of providing advice. Before this, he exerts effort. Now, in order to look at his incentives to provide effort, we have to look at his profits. After some transformations, the profits can be written down in this nice form, which is that the advisor gets a fee, fixed

fee f, plus the retainer and return of profits. What is this term? This is the term that the profits the advisor would get when he would always recommend product A, which is Ta, extracted over the negative, plus the commissions, small Ta, minus the handling costs, k, minus the special costs coming from "rho" when you overstate it, you have to buy product A and the customer follows.

But this is what the advisor would realize when only recommending product A, but then by exerting effort, he can reduce basically the likelihood of a mismatch, and therefore save on the costs "rho".

This, or his incentive to exert effort depends, of course, crucially on his subsequent recommendation decision on q*. For instance, if the commission on product A is so high that subsequently he always recommends product A, of course he has no incentives to exert effort, because the information he gets would be of no use for his decision later on.

So, as an extreme example, effort is always zero when you later on recommend product A or product B with probability one.

So, this is what's happening at the advice stage. After acquisition, after information acquisition, and then potentially biased advice.

Let's now look at equilibrium. First, we are serving customers that are naive. What does it mean that they are naive? We've got a very simple formulation of this, and we talk about this in detail in the paper. We think that naive customers basically ignore the commissions, and they ignore the role of commissions. And therefore, of course, think that the advisor will always give unbiased advice and they expect the advisor to choose an unbiased but-off of q* of one-half.

And what do they think about the effort level? Well, they look through, they think through the advisor's optimal advice program, and this is what their prediction of effort is then, which is the first order condition when the advisor will apply one-half as the unbiased cut-off.

Now, that's what customers beliefs are about q*, which is unbiased advice, and about the effort level that the advisor will therefore exert.

The key thing here, of course, and it's stated in the paper, is the contract design problem with the customer, which is, and I'm just hinting at this, covered by two constraints. The main constraint, of course, will the customer be ready to pay the up front fee, f, and interim constraint, which is when you follow

advice, in particular on choosing product A.

This is the contracting with the customer. One side remark on this, we are using, of course, there's also equal to what the agency is contracting, and just one side remark on this, we are using in this paper a fixed transfer. Of course, with a fixed transfer, we can abstract from the agency problem. So, with the incentives of the advisor and the seller perfectly aligned, we can really think of these two firms as acting like vertically integrated firms.

What's the equilibrium with the customer? The equilibrium of naive customers has it that f is equal to zero. There will be no up-front transfer for advice. What's the intuition for this? Consider, or suppose that f is strictly positive and consider now a changing contract as follows: We consider a decrease in f, a decrease in the up-front fee, and an increase in the product price, so as to still keep the naive customer on his participation constraint. Then the term which I have put down here gives you the marginal change in profits, and it has a way to interpret it, which is it's just the difference between the true probability that product A will be sold, and the perceived probability of the customer, because the customer does not anticipate that the higher price will lead to higher commissions

and, therefore, will lead to a higher likelihood of being recommended product A and, therefore, having to pay this higher price.

In the end, it's optimal for the firm to go all the way and set f equal to zero, and basically let the customer pay for advice, overpay for advice, through basically a higher product price, which is then passed on one to one into higher commissions to align the incentives of the advisor and the firm.

And, of course, the policy implication is then immediate, which is suppose you were banned from paying naive customers commissions, okay, that means here that, for instance, you would require that firms can only pay a commission to cover the handling costs of the intermediary, then this would lead, of course, to unbiased advice, q* equal to one-half. Naive customers would have the right expectations. We can also show that social efficiency is higher.

This is what happens with naive customers and their rational policy intervention. What are wary customers? Wary customers have rational expectations about the commission. Even when they cannot observe the commission directly, they see the product price and can infer that basically a higher product price, of course, will allow higher commissions to align incentives in the

vertical chain.

An immediate implication of this rational expectations is that firms are residual claimants, and therefore have incentives to choose the contract to maximize ex-ante surplus. The outcome is second best efficient.

But the key line here is that still that the outcome may entail biased advice. It may entail a commission which is strictly about handling costs, because the total quality of advice in our model is not only affected by q*, whether the advice is biased or not, but also by the effort that goes into the information and precision.

And because you've got only one instrument, in order to stimulate more effort, it may be optimal basically to accept biased advice. This brings home a point that we have also remarked and addressed in our paper on (mis)selling that we should not forget that commissions may serve many purposes. By regulating you may cause inefficiency because you do not take into account that commissions serve multiple purposes in the model task environment.

The last thing and the last slide I'm having here is on disclosure policy. There is nothing formal I can say about this in this paper. We say more formal

things in other papers. Potentially, one could imagine that giving our naive customers a general health warning could act as an eye-opener. It makes naive customers worry. But this is something, of course, a proposition that would have to be confirmed say in positive evidence, and we are going to run it through in Frankfurt particular regarding this, but if this test holds true, then the implication is that such a simple health warning would lead to higher customer surplus and higher cost efficiency. And this is also important, in our model, firms would not have incentives to provide such a health warning because they're making more profits with the naive customers.

And importantly, potentially at this new model I would suggest a general health warning would be sufficient, so you wouldn't have to disclose exactly what the commissions are, because in another paper of ours, basically we warn against disclosing commissions in details. This is not because of information, but because then disclosing the level may lead to inefficient outcome, because firms will then be tempted not to use commissions, even when the purpose of commissions is, say, to incentivize information acquisition or to incentivize customer acquisition, et cetera.

So, we are now offering the potential instead of banning commission, another policy could be a general health warning. Thank you very much.

DR. YOELI: I'm going to restrict questions in the interest of time. Next to discuss is Heski Bar-Isaac, Professor of Economics at NYU Stern School of Business.

DR. BAR-ISAAC: So, thanks for a great presentation. So, a few interesting things to note about this, which are going to affect my discussion, actually. So, this is a paper that's effectively about how consumers are going to mistrust advice or not mistrust but not trust it enough. Potentially you think of agents and bait and switch and so on and so forth. We had one bait and switch already in terms of the presenters, and the title changed from what I had as well.

And the title changing is actually something substantive, because I'm not an expert in this area. I mean, I think that in some sense, the only experts are Roman and Marco, because people just haven't thought hard about advice as far as I can make out.

So, when I was thinking about advice, and this kind of exercise problems, rather than thinking of these financial applications that they have in mind, I have

personal reasons to think hard about real estate agents.

So, we have this discussion, and this kind of came up with Aviv yesterday. People just don't trust real estate agents. So, I think the naivete regarding real estate agents is that they believe that the sole incentives of the real estate agents, and I am getting nervous because everything I'm saying is being typed, and so the sole purpose is to exploit their naivete or whatever.

So, I know a real estate agent very well who works in rentals and her fees get paid by the buildings that she's renting out. She's getting a commission in the way that's described and the naivete that she faces is that many of the clients that she's dealing with think that her sole purpose in life is to try to exploit them in some way, shape or form.

So, it's nice, first of all, to kind of lay out that these agents are actually doing something that creates some social surplus. Maybe we as economists understand this, but I think the kind of popular perception is a little more skeptical about it.

So, let's just sort of step back and I mean, what is it that these agents are actually doing? Again, what they're doing is, in the model, and I think also in reality, is that they're working very hard. They're

working very hard to find out if the product is well suited to the customer, and once they've kind of figured out, this is your apartment, they convince you that this really is an apartment that suits you well.

Now, in real estate, of course, there's going to be more action by the consumer. After the fact, you've seen the apartment, you're claiming it's fantastic for me. Well, I can see that I have a view of a brick wall, I'm not so convinced. So, there's some information acquisition after the fact as well that potentially can influence this.

So, these agents are working extremely hard to figure out how well suited the product is for the consumer. Once they've gathered their information, they're making some recommendation.

Trade-offs involved here are going to be both in the recommendation and in the information gathering. So, when I'm making a recommendation of product A versus product B, I'm thinking about two things: I'm thinking about, am I going to get a referral from this client, are they going to pass on to other guys that they are really happy with what they have after the fact, are they going to sue me for recommending the wrong thing? So, this is the reduced form reputation costs that they have in the model.

And the other thing they're thinking about is, well, if I recommend this thing where the building is giving me a massive kick-back, I would kind of like to have that flat-screen TV that they are promising me if I send consumers their way.

So, there's this question about the recommendation. Then stepping back from that, there's a question for their incentives for appropriate information-gathering. If I know that I'm going to send them to the building where I get the flat-screen TV, why bother gathering the information to figure out what's going to suit them or not.

So, my decisions of what I'm going to recommend are the things that are incentivizing how valuable it is for me to gather information. If I'm not going to use that information very much or how much I am going to use it is going to affect what I am going to do.

So, what's moving around incentives in this model? Well, consumers have no ability to influence the agents' incentives at all, in the way the model is set up. They pay a flat fee, that's the sunk cost, it's not going to change the way the people behave.

In principle, and I think they treat this reputation cost as something very reduced form and Roman sort of explicitly stated, we're not going to base

policy on this. Actually, this is a very interesting place on which you can base policy.

Teach the consumers to complain after the fact, show them what they could have gotten, have them suing these guys. Potentially, I think there's a lot of scope for policy coming out of thinking a little bit harder about this reputation cost.

Sellers, instead, have only one tool to influence the incentives for the agents. Their tool is this commission. But this tool is a pretty blunt tool because it's doing two things at the same time, it's affecting both the way you recommend, and the way that you're going to gather the information.

So, we've got one policy instrument that's affecting two different incentives, so it's not surprising that we're not going to get to first best here, we're going to be thinking about second best and, the second best does strange things.

If all the consumers are sophisticated we get the second best outcome; surplus is just fully transferred through fixed fees and we're going to maximize the surplus by setting the commission appropriately.

If consumers here in this model are very naive, they're going to believe the agents are going to

recommend completely altruistically. That means that they're going to be more optimistic than they should be about a recommended product and so you would like to kind of push these recommended products harder.

How do you push these recommended products harder? You give a higher commission. The higher commission is going to be associated with a higher price. How do you have that higher price and still keep the consumer interested? Well, you're going to have to reduce the fixed fee down.

So, the fixed fee falls down to as low as it can, which is zero, and we get this higher price. So, these naive consumers end up suffering for the naivete. They end up buying things that are inappropriate. Prices are higher than they would otherwise be, fees are zero.

So, there's plenty that you might think about doing and data in some of the three or four papers out there already or the five or six that are yet to come, some of these things will happen, but it's interesting to think about competition among the firms. Also to think about competition among the agents.

I've highlighted that trying to kind of unpack what these reputation costs are or liability costs are has some implications. So, I mean, is there a

possibility for these agents to make weak recommendations or stronger recommendations that are going to affect how severely they suffer from reputation? Potentially this gives another policy tool for policy makers to think about.

I have an interest in consumer information gathering, and in here, there's no role for consumer information-gathering as a substitute for the agents or even as a complement to what the agent does.

So, if I'm thinking about the kind of real estate example, what happens is they narrow down the choice set for you. I mean, for these financial products, I think the push comes as much from the fact that people are overwhelmed. I mean, I'm thinking of these naive consumers that may be overwhelmed, there's 20,000 different products there, it's hard for me to keep all of them in mind. My financial advisor comes to me and says of these 20,000 things, just think about these 10 or 15, or think about this one and think about things that are similar to this one.

Sometimes I think what they're discussing is narrowing the choice set and then these consumers are thinking harder about that. How does that affect things, does that make much difference? You know, I don't know.

So, similarly with respect to naivete, I'm a little bit more cautious than Roman coming out of what I've seen of real estate agents, where I think the naivete almost goes the other way. I think that people are too suspicious about their real estate agents. They don't take them seriously; they're quite happy to leave them hanging on a Saturday afternoon and ruin their weekends and such, and they don't trust them.

So, maybe we have to be a little bit application-specific in interpreting this, is it really that they believe zero or are they really misconstruing the levels of these things?

Let me just sort of finish up by reiterating that I enjoyed the paper very much. There's really surprisingly little theory on this.

I think just sort of laying out some basic things, being clear about what it is that these agents are doing, thinking hard about their incentives and what's driving them around is very useful, and the clear, simple way that Roman and Marco did in this paper is something that I enjoyed very much.

DR. YOELI: Up next, in about five minutes, we've got Devin Pope, who is an assistant professor of operations and information manager at Wharton.

DR. POPE: So, I'm very excited to be here. So,

I want everyone to imagine for just a second that you're a loan officer and you're thinking about making a loan to somebody, and you've got their credit score, you've got a bunch of information about them, but then you also happen to have a picture of them and a brief statement that they have made for why they want the loan, all right?

So, for example, imagine a couple of the following. So, imagine you have a nice family, they give you a picture of themselves, there's a Christmas tree in the background, a couple of cute kids. They say they need a car. Or here they put a picture of their baby and they need money for medical expenses. Here's a good-looking guy in a nice white shirt that's wanting to become debt free. Here's a dog, and they want to take a trip. So, here's a girl who needs some school supplies. Here's someone who needs their roof fixed.

So, the question that comes with these pictures, well, how are people going to use this information, perhaps, if you were really a loan officer that could use this information, how would you use this in your decision-making?

So, in this paper, what we're going to use is we're going to use a market called prosper.com where this is actually occurring. Our two primary questions

are, well, how do our lenders use this type of information? And then secondly is are they using it efficiently? Are they able to glean information from that picture and the words to actually make more efficient loans?

So, the key advantages of this situation is where it can actually have all of the information that lenders had available, so there's going to be very little unobserved heterogeneity in this context, and we are also going to be able to see how likely people were to actually repay these loans.

So, this is kind of a unique situation where we can measure both, say, discrimination, or how people use these various characteristics, but also whether they're doing so efficiently, so whether it's taste based or statistical discrimination or something else.

I am going to talk about Prosper for those that don't know this marketplace, and then I am going to do two things, first test if there the disparate treatment of the people, and then secondly test if this is efficient disparate treatment, so is it statistical or taste-based, and then conclude.

So, prosper.com is this online platform, where if you're someone who wants some money, you can go there and create a loan listing, which aside from giving your

credit and other information about your employment history to people, you can also post a picture and a particular phrase.

These loans, if you get one, are a three-year, uncollateralized loan, and then what happens is individual lenders can get on this platform, and fund a portion of your loan. So, they can give a minimum of $50, but they can give as much as they want of, say, a $5,000 or a $10,000 loan, and then if enough people bid on your loan, then it becomes funded, as soon as enough people have made it so that the total amount is funded.

And then once it's all the way funded, people can continue to bid down the interest rate. So, you set an initial interest rate, if enough people do it, then it starts to bid down the interest rate.

And, you know, Prosper charges some amount for each loan, and if it's defaulted, then they report it to the credit agency, that's the only incentive for people not to default.

So, here is an example, a lady wanting to pay off her high-interest debt, there's a picture of her, and she's asking for $5,500 at 14.85 percent interest, and Prosper gives their credit grade, and some additional information about the person as well.

Prosper started in about 2006, and as you can

see, there's been a large increase in the number of listings that have been put on this platform. The funded listings have been rising a bit slower, but overall, they funded about 24,000 loans worth $154 million.

So, what we're going to do, if you see those green lines, we went in and we grabbed data for a one-year period. There are about 110,000 loan listings, about 10 percent of which ended up being funded, and we have all the credit information and everything that lenders were able to see about these people, and in addition, we have the photo and the little saying that they wrote down of why they wanted the loan, and we had undergrads go in and systematically code up this information about the photo, so we know quite a bit about this photo that they took.

So, here are the summary statistics. So, most of these people are low credit grade, you can see that 54 percent are coming from that 520 to 560 credit score category. On average, people request about $7,000 for their loan, and they're paying anywhere from 17 to 20 percent interest rates. You know, and this is important to note, about 46 percent of people post a picture. So, a large group just does not post a picture. So, this is worth noting as well.

Why do people want these loans? Most people indicate it's to consolidate or pay down debt, but there's business and entrepreneurship activities, and a host of other things as well.

So, for the sample that have pictures, so of the about 50,000 loan listings that have pictures, you can see that most of them, 65 percent had an adult or several adults in the picture. Ten percent had just kids. Some of them just had dogs or buildings or animals or other things. There's a good mix of genders. Racially, about 67 percent are white, 20 percent black, and then some other more minority races as well.

You can look at age and things, I'm not going to have time to go through all of them.

So, our estimation strategy is basically following the literature on the observational data where we control for as much as we can and then see if these race variables and other variables continue to have predictive power for who gets funded. The nice thing relative to some of this literature is since everything is online, there's no interaction between the lenders and borrowers. We actually have all of the information that was available to lenders, and assuming that we can control properly the regressions, we're going to be able to get these unbiased estimates.

And then the nice thing about the observational data relative to what the auditor feels we can track ex-post performance to see if these were efficient decisions.

So, here are just the raw data for race. So, you can see that for each credit grade, this is the fraction of loan listings that were funded. So, you do see that whites are more likely to be funded than blacks across all credit grades. You see that females, it looks like they're perhaps a bit more likely, but it's not as clear-cut. It seems as if the older people are less likely to be funded.

But let me show you now a regression analysis that controls for all credit information in a very flexible way and we do a propensity score to make sure we're doing a good job of controlling for things. So, here are the key effects. You find that females, if it's a single female in a photo where the base group is a single male in the photo, single females are funded 1.1 percentage points more often than these single males. The largest effects we find are with race, and in particularly the black loan listings.

So, blacks are 2.4 percentage points less likely to be funded than a picture of someone that had a white person in the picture, after you controlled for credit

information and all else.

So, that's kind of the main finding. You can find a little bit on age and whether they looked happy, overweight, a variety of different measures that we had coded up. But I'm going to focus mostly on the race result, since those are the largest.

We can see that as soon as you kind of controlled for the basic credit information, you could include more stuff and you can start including super flexible forms and fixed effects and a variety of things and the co-efficient on black stays about the same. Black loan listings are about two to two and a half percentage points less likely to be funded than their equivalent white loan listings.

So, and then you can say, okay, well what happens to the interest rate that they have to pay? So, they're less likely to be funded, so those that do get funded, you find that blacks are paying on average 0.6 percentage points higher of an interest rate, once they get funded. So, not only are they less likely to be funded, but these loans aren't bid down as much either, and so they're paying a higher payment as well.

I'm kind of screaming through things a bit, but we find large evidence of what people would call discrimination, disparate treatment across races. I

think the natural question is, oh, so are they doing this because blacks really are more likely to default on these loans? So, having a picture of a person that's black is there a signal that perhaps they really will default, even when you control for credit information?

And the idea is that if it's accurate statistical discrimination, then we should see identical net returns for the blacks and the white loans. Whereas if it's taste-based discrimination, then we're going to see different net returns in these loans.

And, so, here are kind of the summary stats up until now. These are all three-year loans, and since this is a relatively new website, we don't have the full information what's happened all the way through, but for a lot of loans, we have about two-thirds of the loan repayment. We go by the standard literature that says if it's more than three months late, we're going to assume that's a default.

So, you can see that about 30 percent of the loans for blacks have defaulted, or four-plus months late, whereas only 15 percent of the white and Asian loans have defaulted. So, this has suggested that maybe black loans are defaulting at a higher rate, but we know that they have worse credit scores from what I showed you and that they are paying a higher interest rate, so

that's reasonable.

So, the question is does the extra that they have to pay in interest rate make up for the fact that they are defaulting more? And I think this is kind of the key graph. I am just showing a loess plot by black and white, so this is the final loan interest rate that they're actually paying, plotted against the fraction of loans that default.

So, here you're funding, say, a white loan at 18 percent, and this is the fraction of default, or if you're funding a black loan at 18 percent, this is the fraction of the default.

So, this is opposite of what would be considered taste-based consideration of blacks, which is talked about a lot, and in fact, blacks are defaulting at a higher rate than you would have expected given how much they're being charged in this marketplace.

Another way to think about this is if you look at the start, this is a loan age of zero, so at the beginning, the cumulative APR that people are making on these loans, it starts higher for blacks, right, because they're paying a higher percentage rate to begin with, and then we say as soon as it's been in default, we're going to say it's never going to be paid off again.

What happens is you see these lines start to

quickly cross and the black loans end up being not as profitable. So, you can look at this now with a smaller sample, but across a longer time period you're seeing the same thing.

Notice people aren't doing very well in this marketplace, right, even the white loans, you're only getting about a four percent return after about a year, so this isn't a great place to invest in, at this time, at least.

You can also show this with regression analysis.

Let me just wrap up with a couple of conclusions. I think one thing that's clear is that skin color isn't a causal factor of default, almost surely people would argue that what's happening here is that there are factors related to being black that perhaps could be correlated with default. So, maybe education levels that aren't represented in the data set or neighborhood effects or support networks, or it could be labor market discrimination itself, which is causing them to have to default more, because they're being discriminated and can't get a job, right?

So, there's a lot of reasons why we might expect blacks to be defaulting at a higher rate. Given that we're finding lower net returns amongst blacks, what's our conclusion? Well, we're finding statistical

discrimination against blacks, remember, because they are being discriminated, there's disparate treatment, but then we have to argue that this is taste-based discrimination against whites, because the statistical discrimination doesn't go far enough.

So, that's kind of a strange conclusion to make, right, do we think that there's taste-based discrimination against whites? Well, an alternative explanation is that people just aren't very good at knowing what's the propensity of default for these various groups.

So, lenders are unclear that being black is apparently in this market a very strong signal for defaulting. So, rather than being taste-based discrimination against whites, it could just be inaccurate beliefs by lenders.

So, and you might expect this to disappear in the long run in this market, as consumers, or as lenders perhaps become a bit more savvy about this, but it's been around for three or four years and we haven't seen it going away yet.

So, what are the take-home messages of this? I think one of the most important things that comes out of this is a lot of times we try to look for taste-based versus statistical discrimination, and occasionally,

people find that there's taste-based discrimination against blacks, for example, and we conclude, all right, well there's taste-based discrimination against blacks. Well, in this case, we find taste-based discrimination against whites, and we say, well, let's hold on here, that can't be right. Let's interpret the evidence as inaccurate beliefs.

Well, it could be that when there's taste-based discrimination against blacks, it's also inaccurate beliefs, right? So, we have to be a bit careful about how we're using the various tests that we do when we're trying to test for these different discriminations when inaccurate beliefs could be going on.

But overall, we are finding discrimination and disparate treatment in this market, but it doesn't seem to be quite the type of disparate treatment that we're used to, when we think about race.

Thank you.

DR. YOELI: Any questions for Devin? You want to come back up? You have a question.

DR. POPE: I'm sorry, I apologize.

DR. ROWE: Hi, Brian Rowe, FTC.

So, you have this claim that accurate statistical discrimination by the individual lenders will equalize the net returns, but if that were

literally true, then you could simply compare the net returns across the races and be done. Instead what you've done, which I think is more reasonable, is pool the data across lots of different options, right, with different borrowers, and I think that makes more sense, because there's a result in a couple of Nicola Persico's papers about how in these types of matching markets, what you need to equalize are the returns across all the prospective borrowers when matched with all of the prospective lenders.

So, I don't know if that happens here, it seems like it may not. So, if true, that's a case that makes it even stronger for doing what you did here.

DR. POPE: Good, thanks. I think I agree with you. So, I think part of the issue that you're bringing up is you have to be careful when comparing net returns. So, say, in a real estate market, let's see if blacks default at a higher rate than whites on homes or something. You have to compare marginal rate, you can't compare distribution, that's one of the things brought up in this literature.

The nice thing here is because it's a bid down process in the interest rate, we're able to actually compare marginal versus marginal, and so it does work out, which makes for a nice test. Thanks for the

comment.

DR. GENAKOS: Christo Genakos. Just a clarification, do you control for the size of the loan that they were asking?

DR. POPE: Yes.

DR. GENAKOS: Right. A suggestion, you mentioned that some of the websites didn't provide their photographs, so you could compare blacks with photographs versus black or white with or without photographs, and I would like to control there for whether their name clearly suggests that this is a black person or not a black person.

DR. POPE: Good, great question. So, there are no names, it's all kind of anonymous, although there's pictures, right, but there are no names. So, we can compare, say, what's the value of having a certain photograph versus not even posting a photograph. So, this is a bit of a weird market, and it's kind of cheap talk, right, to put the photo up there. They don't go and make sure that's you. You can put up any photo you want. We see a lot of people kind of being happy to reveal this information. There is a lot of value of putting a photo up, so the people that don't put up any photos are treated about the same as putting up a photo of a black person is treated.

So, I mean, it suggests that if you're a black person. The best thing you can do is put a picture of a white woman on your loan listing, and the worst thing you can do is put a picture of a black person or no picture at all is what we're finding.

DR. STERN: So, Scott Stern from MIT and Northwestern. So, we didn't hear very much about who's actually lending the money, and it would seem to me, I'm thinking kind of John Gorian's kind of studies of taste-based versus statistical discrimination.

DR. POPE: Yes.

DR. STERN: But what really matters here is not on average are people discriminatory, but what's really happening on the margin, and you can imagine that these three groups of people, there might be black lenders who are actually seeking out black borrowers, you can imagine a whole group of people who have no discrimination at all, and then you can have a margin, some guys on the margin who are engaging in taste-based discrimination.

But it was unclear to me whether or not the regressions that you ran here are going to allow you to do that matching. In particular, the fact that you see lower lending rates, but also higher default rates would suggest something about, oh, is there a subsidy on the

lender's side of the market, and that might suggest either the misinformation, kind of the inaccurate police story you're saying, or for all the people who are lending the money are people for whom, in fact, this goes fairly well.

DR. POPE: Good. These are great questions. So, a couple of points. Unfortunately, we don't know a lot about the lenders in this market, so we don't know the lender's race, for example, which would be a nice way to kind of test an own race type bias.

So, in our paper, we have a model of a representative agent as a lender, but you're right, it's important to think about how this could be heterogenous and it must be that the marginal lender is the one causing these differences.

So, yeah, I think that's absolutely right. So, it's important to kind of think about what exactly is this meaning about the entire population, it's not quite clear, but it does tell us something about kind of what's happening right there at the margin.

DR. YOELI: Last one?

UNIDENTIFIED MALE SPEAKER: So, I have a follow-up question about skill set, that to me it wasn't clear whether this is an investment place or a charity place.

DR. POPE: I didn't answer that very well.

UNIDENTIFIED MALE SPEAKER: I would think that the data you got, every investor, not exhaustively, but they could get some statistics on what the default rates are.

DR. POPE: Yeah. So, it's not clear whether or not this could be some charitable giving mixed with kind of people wanting to make money. Most people suggest that they are in here to try to make investments, and interestingly enough, all the other information kind of matches very well with people being rational about their investment.

So, for example, higher credit score people are much more likely to be funded and get a lot better loans. So, everything works as if people are trying to pick the best bets. So, it looks like people are not being completely charitable, right? So, the low credit people don't get all the loans at the best prices. That's definitely true. But it's kind of a weird marketplace.

DR. YOELI: Thank you. Up to discuss, Patrick McAlvanah of the Federal Trade Commission.

DR. McALVANAH: So, one of the difficulties with the discrimination literature is this finding that if you test for discrimination, but you don't find it, it's

hard to tease out whether there is no discrimination, people aren't discriminating, or whether the market is taking care of it and priced it out.

So, in my mind, this paper has two key methodological advantages: The first, as Devin just said, it's fundamentally a weird, thin and quirky market. This allows us to get a little bit closer to the black box of what's going on with borrowers at the lender level without some of these thick institutional veils.

And the second advantage, as Devin conceded, was the omitted variable bias. So, any time we're testing for discrimination, there's always, it pops up so frequently that when you look at the raw data, it appears that there's discrimination, but then once you start adding controls, the discrimination goes away, or vice versa. The concern is we might accuse someone of discrimination, but a lender or some institution might have had more information than we can observe, and then once you control for that, it would have gone away.

So, the advantage of this prosper.com data is that in his regressions, he's able to use all of the possible information available. He has the exact same information set as the borrowers and the lenders. So, that's another advantage of this data.

What does he find? As he said, when you just look at the raw data, of the loans that have voluntarily posted a picture, then the loans associated with a black picture appear to be getting a raw deal. They are less likely to be funded, and when they are funded, they are funded at the higher interest rates. But when you add in the conditional controls for the credit information, then they're almost getting a good deal. That is, there is statistical discrimination against them, but not as much as would be at optimal.

So, what are my comments? As already suggested, it seems that the decision to post a picture is endogenous, and that's a really important decision here and that's what's driving a lot of things. I know you're bumping up against space limitations for the journal already, but it seems to me that it's an important question to find out, are the characteristics of the loans without a picture the same as those loans that do have the picture?

In your first table, you had just one column that is the summary statistics for all the loans, and then the next column is the summary statistics for loans that are funded, I thought you could pretty easily just add one additional column for the loans with a picture, just to see if they're comparable to the full sample

sets. You say it's not in any of the regressions, but it's in the text, and as you just said, this decision does seem to matter that posting a picture, any picture, leads to three percentage points, relative to about 10 percent sort of loans that are funded? So, that's a big difference.

There's one additional test, and keeping on the same subject, an additional test for accurate statistical discrimination is if people are profit maximizing, well then the net return from loans with any picture should be at least as high as the loans without a picture. Just because there's more information to optimize and to use in your statistical discrimination calculation.

So, you could test the contra positive for this statement, that is if you see loans without pictures are more profitable on average than the loans with the picture, well then that should imply then that people are not profit maximizing or that there's not accurate statistical discrimination going on.

Here's a tidbit from the psychology literature, there is what's known as the identifiable victim effect, and this is the phenomenon whereby individuals are more likely to help an identified victim rather than an unidentified victim. So, if I could say, hey, would you

like to donate some of your experimental winnings to a person randomly selected in this room, you will get a low level of donations, whereas if I specifically say, your winnings will go to Paul, people would be more likely to donate then. Or if you say, hey, would you like to donate to this child who is sick, you will get a lower level of donations, where if you show a face, and you show a picture and you show the name of the sick child, you get a higher level of donations.

I don't want to get too far aside, but in every single presidential election that I have been alive for, what you always hear when the presidential candidate is talking about their health care plan, they always say, well, here's why my health agenda is better, consider the case of Ida Thompson, an 80-year-old grandmother in Plano, Texas, who now has to work a second job or pick up another job because Social Security isn't enough to cover her drug expenses.

They always single out one person as opposed to the statistically more compelling, consider the 20 million people who would be affected by this policy. Presumably, because it allows us to empathize a little bit more, it tugs at our heart strings. This effect actually reverses at the group level, so if you say would you like to donate to these eight sick kids,

generally, you'll get a higher level of donations than if you specifically point out these eight individuals, and actually reverses for in-group versus out-group.

If you go to Israel, and ask an Israeli citizen, would you like to donate to this Israeli soldier who was wounded, you will get a low level, while if you single out, would you like to donate to this Israeli soldier, this specific one who was wounded, you will get a higher level. But if you ask an Israeli citizen would you like to donate to this Palestinian citizen, in general you will get a higher level than if you single out this specific Palestinian soldier who was wounded. So, that reverse at the group level.

This was actually the first thing I thought of when I read this paper, that seems to be what's going on. That is when you see along with a picture, you can identify a little bit more and might be able to put yourself in their shoes, and, in fact, like we said before, if we see the loans with any picture, more likely to fund.

For what it's worth, you don't pick up that effect at the group level, it's insignificant. Then it would be really hard to test for the in-group/out-group bias because you would need that information on the lender characteristics, you don't have that, so I

thought that was interesting.

And then my final point, again, you concluded that there is accurate statistical discrimination against the loans with a black picture in them, plus some taste-based discrimination in favor of blacks or for whites, or this boils down to insufficient statistical discrimination or inadequate adjustment to it, and that would have to imply bias belief surviving at the market level.

On this issue of when can an individual level irrational or an individual-level bias survive to the market level, I am a behavioral economist so I get attacked on this eight days a week, and the determinant for a bias surviving to the market level is whether the actions of rational or irrational types, strategic compliments or strategic substitutes, and then how thick is the market.

So, if the actions of the rational and irrational types are strategic compliments, well then even as n goes to infinity, it's not a no-brainer that this bias will not disappear, because you get rational types wanting to mimic the irrational types.

If they are strategic substitutes, then that's when you can make the appeal to, okay, they're probably the irrational types are going to be making less profit,

they might be priced out of the market. I'm more sympathetic to that at the firm level, because if the firm isn't profit maximizing, it's probably going to go bankrupt. It probably has a little bit less appeal at the consumer level; it's never been instantly clear to me that if the person is not maximizing his utility, they're instantly going to disappear.

What's it look like in here? If you are not perfectly engaging in statistical determination, you are not making as much profit as you could be. So, I am sympathetic to that. It's hard for me to imagine, though, anybody using prosper.com as their primary source of income. So, I'm willing to admit that they probably are strategic substitutes, but the evolutionary selection pressure probably isn't that strong.

However, on this issue of how thick is the market, it is not instantly clear to me that this is a very thick market, and precisely for the reason that it's an interesting study in the first place for discrimination is that it's kind of a thin, quirky market. That's sort of what makes me not sympathetic to the appeal that, okay, this is a market that is going to price out the bias beliefs.

So, that's just my two cents. I think that's it.

DR. YOELI: Thanks, Patrick, and thanks to all the speakers and discussants, we are running a little bit late. Five-minute break, so be back here at 25 after.

(Applause.)

(Whereupon, there was a recess in the proceedings.)

DR. STERN: Okay, let's get started with this I guess penultimate piece of what I think has been a really great conference over the last two days, and what we're going to do here is kind of open things a little more loosie-goosie, free form in terms of trying to figure out kind of a broader set of issues and really kind of raise up some ideas about, in particular, what the new administration's innovation policy agenda is, and also what it should be.

So, what I'm going to do is I'm going to vaguely dispense a little bit here with doing too much introduction, because we have a short amount of time, and basically what we're going to do is we're going to hear from three very distinguished panelists, and then we're going to kind of mix it up between them and hopefully some questions from the audience.

So, I'm going to start with Joe Farrell, who obviously needs no introduction in this room, and so why

don't we start with Joe.

DR. FARRELL: Thanks, Scott.

So, a couple of disclaimers. First, as you know, I am speaking for myself, not for the Commission or any commissioner. Secondly, some of you may be less aware, even if somebody were to speak for the Commission, they wouldn't be speaking for the Administration, the Commission is an independent agency and not an arm of the Administration.

So, what about innovation policy, as far as I think about it? I think about innovation, to the extent that I think about innovation policy, as a branch, an application of competition policy, for the most part, although I will have one or two other things to say in a few minutes.

I think the issues in innovation-oriented competition policy are actually to a great extent the ones teed up by Scott yesterday, when he talked about the way in which market incumbency might relate to either transacting or doing your own research and development, and what are the incentives for others to come and do that.

I wrote this in the form of a slogan, incentives matter for innovation as they do for other things, but I want to make two qualifications to that. The first is

that incentives are not just for incumbents. So, we're used in the Washington economics world to hearing about incentives for innovation largely, not entirely, as an argument for why one shouldn't mess with the rights of established incumbents to collect the fruits of their legitimate monopoly. I think there's a lot to that.

Of course, from a competition policy point of view, I do like to stress that the fruits of the legitimate monopoly presumably do not include the fruits of what's called an antitrust monopolization, because in what perhaps sometimes risks being a slightly circular way, monopolization is defined as going beyond the legitimate monopoly through innovation.

Nevertheless, although it may sound slightly circular and sometimes be slightly circular, I think that is well worth considering, because although the incentives for incumbents argument is a legitimate argument for not going after incumbents past a certain point, it's not necessarily a very good argument for going soft on monopolization concerns.

So, incentives matter, but they're not just for incumbents, and formally, in the economics literature, where we see this coming out is primarily, I think, in the idea of complementary or follow-on or cumulative innovation. When you have strong complementarities, the

arithmetic of incremental contributions is that you can't pay everybody their incremental contributions, so there is an inevitable tension between getting the incentives right for the provider of component A and getting the incentives right for the provider of component B. You can't do it at least with budget balance and so that means it can't be done in a decentralized market oriented way, of course you could do it with a very interventionist taxed and subsidized policy.

So, it seems to me that there's a fairly common fact pattern exemplified most prominently these days in the area of broadband Internet access and in neutrality, where in component A, you have a relatively stable, well-identified firm or set of firms, and in component B, you have something that might be well described as an eco system, rather than a set of identifiable players. There is a real question then of how do you decide, and what are the gravitational forces pushing you to stress the incentives for one of those layers versus the other.

I think there is a risk when people in this town talk about innovation incentives, of letting that mean, by default, incentives for the identified lobbying, clear-cut, already in existence, can point to investment expenditures players as distinct from the garage-based,

not necessarily identifiable, not even necessarily in existence players.

I think some degree of resistance to that tendency, if I'm right in identifying it as a tendency, is what underlies the innovation part of net neutrality policy.

That brings me to my second qualification to the incentives matter slogan, incentives matter, and my first qualification was not just for incumbents, also for the guys in the garage, but it's also true that in many, although not all areas of innovation, while incentives matter, incentives are not the only thing. In particular, in an economic model, it's typically true that if anybody is able to do x and you give them the right incentives to do x, things work out well.

In the real world, I think, and certainly in the world of ideas and imaginative innovation, you can give people incentives to come up with bright, new ideas, and they very often don't. You can take away people's incentives to come up with bright, new ideas, and they sometimes will.

So, if we look at Silicon Valley, for example, any competent economist looking at the organization of Silicon Valley could tell you the weakness in California law of noncompete agreements, and the tendency of ideas

and innovation assets to walk out the door each evening, would ensure that Silicon Valley could never be the home of a significant amount of innovation. What that lesson suggests, and I believe it's often true, is that opportunities for large numbers of players can matter just as much as getting the incentives right for small, well-defined sets of players.

So, those are the two qualifications that I think one should keep in mind, or that I try to keep in mind, to the perfectly legitimate, but somewhat conventional economic message that, of course, incentives do matter. Thank you.

DR. STERN: Great. Thanks, Joe.

I think very provocative comments, particularly what we will come back to on this issue of sometimes you put a lot of incentives out there and that's how economists tend to think about innovation, but very often the creative process is a bit tricky and thinking about how we incorporate that into policy is quite challenging.

Joe yesterday when he introduced me talked about learning to not try to do too many things and time management and the whole thing, and our next speaker, I've always believed that Bob Litan has, there must be two of him, because there's one Bob Litan who has a law

degree and has had a variety of positions in government, both at the OMB, at the DOJ, so on and so forth, and then there's another Bob Litan who has a Ph.D. in economics, and among other things, was the head of economic studies at the Brookings Institution, and both of those Bob Litans, apparently, are currently the vice president for research and policy at the Kauffman Foundation and continue to be a senior fellow in economic studies at Brookings. So, without further ado, I'm looking forward to hearing what Bob has to say.

DR. LITAN: Thank you, Scott. Both of those Bob Litans are getting tired, and so I'm winding down.

I'm going to make three points. The subject here is the administration's innovation agenda. I'm going to grossly oversimplify it, but basically it boils down to spend more money on basic R&D, and focus it on three basic areas. On clean tech, health care IT, and the third one is probably education, right? I think those are the big three.

So, coming from Kauffman, I mean, by the way, I just want to let you know, I don't have any objection to spending more money should you find it, and certainly you can make an argument, the classic public good argument for spending more money on basically all three, but coming from Kauffman, what we care about is the

diffusion of the knowledge that's generated by all this, into the real world, and we know the way you get diffusion is through good commercialization. In particular, the entities that are most likely to commercialize are new ones, rather than existing enterprises.

And it's not just the litany of radical innovations, from the car to the airplane, air conditioning, all the examples that we like to cite, all of those having been done by entrepreneurs. It's not just those that underscore the importance of new businesses, but we've also published some new research that's on our new website, we had the Wall Street Journal write about this about ten days ago, that documents that new firms, it's not small firms, but new firms have accounted since 1980 for all of the net new jobs created in the United States.

So, you want entrepreneurs, both because they're radical innovators and because they're job creators, okay?

Keeping those ideas in mind, and assuming that we're going to spend more money on R&D, how can we make sure that this stuff gets translated and commercialized, especially by new firms?

I'm going to give you two generic ideas, and

then one specific one on health care, which I know is dangerous in the current environment. The first idea: Change immigration policy. I mean, here I may be speaking to the converted, but it's not true. There is not conventional political wisdom, but clearly, we know from the studies that we've commissioned, immigrants are disproportionately successful in forming high-tech, successful companies. Work from both Duke and Harvard documents that 25 percent of all successful high-tech companies of the last decade were founded or co-founded by immigrants, that's consistent with work done by Joe's colleague at Berkeley, Anneliese Boghossian, and you certainly know this, of course, if you live in California.

So, ideas like giving graduates, immigrant graduates of American universities who have STEM degrees, giving them green cards with their diplomas is a great idea. And if you can't swallow that politically, then let's think creatively maybe as a fallback option what I call a job creators visa and let's have at least some renewable visa for immigrant entrepreneurs who are hiring at least one or more individuals.

I know there are technicalities regarding such a visa, but clearly immigration policy and keeping smart,

educated people here is good for this country. It doesn't displace jobs, it creates jobs. So, that's point one, and that's by the way, a way of ensuring that a lot of this high-tech stuff that's going to be funded is going to get commercialized. That's point one.

Point two: Speaking of competition, there is one area of university activity that actually is still a monopoly. If you are a university faculty member and you are a rocket scientist, or less, and you come up with a brand new idea, first, your university owns you, and owns your idea. If you want to commercialize it under existing law at virtually all universities in America, I think without exception, you are obligated to go to the university's technology licensing office and disclose it. And, you are required to use the university's TLO to license or commercialize the technology.

You cannot go to your own lawyer or you cannot go to the lawyer of another university's TLO who may actually know a lot more about your technology than your home university, which is probably stretched thin, is looking for the homerun, is looking for the next Google, may well ignore you and put you in the back of the queue.

So, there is a bottleneck, in our view, created

by monopolies, that are artificial, at every major research university in America. So, what we need to do is unleash or get rid of that monopoly. We need to basically allow faculty members to choose how to license their technology. This would not change the distribution of royalties between the faculty member and the university, but it would accelerate the commercialization of the actual technology.

There's one very simple change in federal policy that can make this happen tomorrow. The Federal Government, which hands out billions and billions of dollars, could say, tomorrow, to x university, we're not giving you any money unless you allow your faculty members freedom to choose their technology licensing agenda, period, end of story.

Now, you may want to have a time limited thing, which says, if the individual faculty member's lawyer or their individual does not commercialize within a year or two, then it reverts back to the home university's TLO, but the fundamental principle ought to be that we end the monopoly, one of the few monopolies left, actually, in America. We ought to end this monopoly at each university and allow a free market in technology licensing. That's the next big idea.

The final idea, and I know I'm treading on very

thin ice here talking about health care, but look, everybody is, these days. We know that the key to getting health care reform is bending the cost curve. We also know as economists that the major reason for the unending or unceasing drive towards higher cost health care is technology that is more expensive. In many cases it's good, it makes for less painful procedures and so forth, but in a system where everybody is insured, there are tremendous incentives to develop more CT scans and MRIs, okay?

So, we have basically a system that is biased towards cost-enhancing technology. We need to find ways to give incentives for cost-reducing technology. And I'm not a believer that the government should be in the business of picking winners and losers, and, therefore, I'm a little worried about comparative effectiveness analysis, because while it sounds good in theory, once you introduce comparative effectiveness analysis, you may be stultifying or dampening innovation. Because you're dictating what procedures will get reimbursed, and by definition, that will discourage innovation. That's not what we want in health care.

I think, as economists, getting back to incentives, there's only one thing we know how to do to encourage the right kind of innovation, and that's get

the incentives right. So, we have to find a way to have individuals bear more than the first dollar cost of their insurance, which is not the direction we're going in. Secondly, I think we need to find a way to encourage providers to join integrated networks, like the Mayo Clinic that the President talks about, Paramount Health Care, and so forth. These organizations have been very successful in delivering low cost and very effective health care, and they pay their doctors on a salary basis. In fact, ironically, we have too much entrepreneurship in health care because we have too many doctors who are operating individually and not as cohesive units. We have to figure out a way to change our reimbursement systems so that we encourage more doctors to join these integrated networks, and once we pay for things in a less than fee-for-service way, and pay for outcomes, rather than inputs, then those providers have incentives to use the technology that saves them money, because it will save them money as well in delivering health care.

So, I think a great challenge for the profession is to figure out innovative compensation arrangement for health care providers. I think we need to move away from first dollar insurance, and unfortunately, as I said, that's not where we're going.

I'll end there.

DR. STERN: So, in addition to the administration's innovation policy, we're also going to tackle health care. But I think you're absolutely right, that the challenges around health care innovation are inextricably tied right between this innovation policy piece and the health care policy piece.

Our final speaker, and then we will see what commonalities we can start to draw from these presentations, is Tom Peterson from the National Science Foundation. He is currently the assistant director of the NSF for the engineering directorate. He comes to that position after being dean of engineering at the University of Arizona. We are doing a session on innovation policy and we finally have somebody who could actually do some innovating, so that's good.

Among his many accomplishments on his CV, he is a fellow of the American Institute of Chemical Engineers, an institution I actually studied when I was in graduate school in dusty libraries at Stanford. He has also won the Kenneth T. Whitby award from the American Association For Aerosol Research, and I can almost guarantee you he's the only person in this room to have done so.

DR. PETERSON: I'm trying to decide whether I've

been ambushed by being invited to this meeting or what, so I'm involved in one of those government agencies that hand out billions and billions of dollars to universities, and being an engineer, I --

DR. STERN: Thank you.

DR. PETERSON: Right. I also have a PowerPoint presentation, which I am going to go through very quickly, because I think it illustrates some of the issues that have been brought up.

So, let me just talk a little bit about our role, and this is going to be a fairly narrow perspective, and I will say that up front, from the point of view of the engineering director at the National Science Foundation. Everybody knows, I think, that the primary focus for the NSF is to support basic research in science engineering and science engineering education, and we have no desire within the engineering directorate to move off that. But obviously within engineering as a profession, you have a responsibility and a role, by virtue of the discipline itself, to have interactions with business and ask questions about whether what you're doing has direct application in a commercial environment.

So, as we look at our role in innovation, I have to say, as we've studied this, I have come to the

conclusion that there are a lot of very innovative definitions of innovation, and for our purposes, we focus on a fairly narrow definition, and that is those activities that have been supported by NSF that have led to some quantifiable economic benefit.

It is this question of what is the relationship between the support for basic research and eventual commercialization. Before I get to that, let me also just say in deference to my colleagues at NSF who look at the social and behavioral and economics side, there are obviously other issues with respect to innovation, and a couple of them were mentioned already.

For example, this question of what is the impact of foreign students in engineering, I can tell you, two-thirds of all the graduate students in our programs throughout the country have their undergraduate degrees from an institution outside of the United States. They are critical to our success.

So, one of the issues that this group called the Science of Science, Innovation and Policy that NSF is looking at is the contribution of foreign students. I don't want to in any way minimize those kinds of issues, they're important, but they're not really issues that we as engineers have the capability to address.

In addition to this focus on basic research in

engineering, the engineering directorate really wants to have as part of our portfolio, translation of research. What we simply mean by that are those kinds of activities that are interdisciplinary by nature and involve teams, often teams that include businesses and industries, and with the expectation or the hope that the results have some clear benefit to society.

I'm not going to go through all of this alphabet soup of different programs at NSF, but let me characterize them broadly. We've got programs that support multi-institutional centers. So, this would be multiple universities often in partnership with other entities or other government agencies.

We've got support that focuses directly on partnerships with industry, and here's a list of them. GOALI, for example, this Grant Opportunities For Academic Liaison with Industry supports bringing industry people into the university or university people into industry to appreciate each other's environments, and hopefully make contributions. Then, of course we've got programs that support very specifically individual PIs.

Everybody has seen this chart, the graph of investment between discovery and commercialization, and as you might expect, if you look at where NSF's

investment is, it's pretty far to the left. If you just put a point where the centroid is, for overall NSF investment, it would be pretty far to the left. Even if you look at the more applications-oriented directives, like engineering or computer information science engineering, we would still be pretty far to the left.

But if you look at all these other programs that we describe, we populate this graph in various areas. The science and technology centers, pretty basic, but these industry university cooperative research centers in which the NSF is only a minority partner, really have quite a bit of influence by ideas from industry. A lot of these other programs you would find in various places.

The engineering research programs are rather interesting. I don't know if any of you are familiar with these programs. The program itself has been in existence since 1985. NSF only supports them for ten years, and yet all but about six of those engineering research centers that have been started are still in existence, because they've developed strong partnerships with industry, and have had their support continued long after NSF support has diminished. So, I think they play an important role in translation of research.

I just want to very quickly in about two minutes

give you some examples of where, in spite of these issues that I think were very correctly pointed out about the tech transfer challenges and the monopoly that exists that individual faculty members have to deal directly with only their tech transfer office, we've been pretty successful in taking these various areas of support to commercialization.

A quick disclaimer, kind of like the disclaimers you heard already, we're not claiming sole responsibility for these successes, but in every one of these that I mention, and I'm just going to mention a few, NSF played a clear and definable role.

So, in 1985, Andrew Viterbi and six of his colleagues formed a company called Quality Communications and got an SBIR from NSF for about $300,000 to develop a decoder, which actually turned out to be a critical element in their data transmission via wireless and satellite. You all know now, Qualcomm is a huge company, a $70 to $80 billion company.

Looking at engineering research centers we are supporting right now an ERC looking at synthetic biology at Berkeley, and Jay Keasling and his colleagues developed a path to synthesize and anti-malarial drug called Artemisinin. It's really important to have this process because it occurs naturally in only very small

quantities and is very hard to extract. Substantial commercial and health implications there.

An example from these IUCRCs: We support a program at the University of Arkansas that partners with Sam's Club and they've developed some software that really helped with their inventory and logistics, and have saved them millions of dollars in inventory costs.

Even the science and technology centers, which are very fundamental by their nature, have some really outstanding success stories. In the nineties we supported an SDC at the University of Illinois on magnetic resonance imaging, and, in fact, the principal investigator of that, again, focused very much on the fundamentals, and won the Nobel Prize in MRI.

Even individual awards have some successes, and some great examples, including support for graduate fellowships in separations processes for liquids and gasses, and these have important implications in clean technology.

And Chad Mirkin, who actually sits on the President's Council of Advisors on Science and Technology, is a prolific inventor at Northwestern, and he's developed some techniques for nanolithography and nano fabrication, again with the support of NSF. All just quick examples.

The point I'm trying to make here is in all these examples, university research is key. Often times it's driven by industrial needs, and you'll see faculty who have various abilities to interact from the basic research all the way through interactive in industry. Not all faculty are predisposed to developing business plans, but I think we can play an important role in smoothing this transition from development to commercialization.

So, that's our plan within the engineering directorate at NSF to grow that portfolio of translational research, hopefully to expand the research for industry-driven initiatives, and support some of the educational aspects of innovation and to partner with our colleagues in social, behavioral and economics science on some of these very substantial issues.

DR. STERN: Great, thank you so much. So, let me just get us started and hopefully can open it up to questions for a few minutes. One of the themes that I think came across all three presentations here is both the centrality, but perhaps the kind of hard-to-nail-down positioning of the role of entrepreneurship of various types, particularly technology entrepreneurship in sort of promoting innovation, economic growth, sort of economic dynamism

through competition.

I guess one question I would throw out to you guys reflecting on what each had to say is, if we wanted to promote a higher level of high-impact technology entrepreneurship, the U.S. has been historically quite good at that relative to kind of everyone else, as opposed to health care where we go to other countries to learn what to do. This is one where we're going to have to extend the frontier.

What can we do from a policy perspective that might significantly enhance the rate of kind of high-impact technology entrepreneurship that promotes competition and productivity growth? Maybe that's something that we're not currently doing? Bob?

DR. LITAN: Okay, so apart from championing my TLO competition idea, which I put on the table, all right? Let me tell you just something briefly that we are doing at Kauffman, which we want to try to scale up, and hopefully it could be a model for the rest of the country.

So, here's a factoid. Do you know how many post-docs there are in the United States? Actually, Tom, you can't answer this, because you probably know the answer. How many post-docs? These are people now that are relatively low-wage people that are spending a

number of years trying to enhance their resume that have got STEM degrees and their job or their vision in life is to become assistant professor at Arizona, or fill in the blank?

Do you know how many there are? There are 45,000 of them in America. All right? And their ambition is to become a professor, and our belief at Kauffman is that there are entrepreneurs in that group that don't even know it, and are sitting on maybe not the next Google, but maybe the next Qualcomm, or something else, and they just need somebody to tell them that and pair them with people who can make that happen.

We have just launched a program called Kauffman Labs, where we have run a competitive program of post-docs in the United States, they are paired with star scientist mentors at their home university, people like Bob Langer at MIT, and we picked 13 out of a class of 330 in the first application pool. We expect probably a thousand applicants next year, and we hope to have a lot more than 13, hopefully maybe 50 next year.

We have a great network and we're going to hook these people up with mentors and networks to help get them either entrepreneurial partners and/or money and/or suppliers or God knows what, and turn them into businesses. That doesn't mean they'll all become

entrepreneurs, they'll go back and teach and maybe be a CTO at the firm, but we think we need to basically birth more firms out of that cohort, and they don't know that they're entrepreneurs.

So, we are, and I'll summarize, on a campaign to recruit entrepreneurs out of a potentially high-growth pool. So, if this model works, this should be scalable. That's the idea.

DR. STERN: Tom?

DR. PETERSON: I'm going to ask to go next simply because Bob and I did not meet before this meeting.

DR. LITAN: But we should continue talking.

DR. PETERSON: Absolutely we should, and he may not even be aware of what we're already doing together. We started this year at the foundation, and within the engineering directorate, an innovation fellowship program, and in fact, the focus is precisely on this cohort.

Let me just make one caveat, and I think it's an important caveat. When we talk about the number of post-docs, it's also important to define in what areas they represent. It is much more common in the sciences, for example, to go through the post-doc route before you go into a faculty position. Not as common in

engineering, although it's becoming more so.

But this is precisely an area that we, too, think is important to support, so we are partnering with industry to provide fellowships for 40 post-docs. We are putting up about two-thirds of the money. And industry the other third, and the hope is -- and this is where I was hoping we might have known a little bit more about this -- that we have been in conversation with Kauffman to participate in this two-week boot camp that you have for your fellows. So that we're going to send our fellows to precisely that boot camp. It is to take them from this point where they have developed a tremendous technical knowledge, but don't have a clue about what to do with it from an entrepreneurial standpoint, and I think it's a great idea.

DR. STERN: Joe?

DR. FARRELL: My grandmother had three sisters, and two of them became psychiatrists. Back in those days, it was not, I think, particularly easy for women to do that, and the same privileged background that enabled them to break through those barriers also tempted them to give up when they discovered that in order to become psychiatrists and psycho-analysts they had to learn, as they used to put it bitterly in their later years, all about the bones of the foot, because

you had to have a medical degree back then to be a psycho-analyst.

I think it's a bad idea to have systems where people can't specialize, and so although there's certainly nothing wrong with creating and opening up opportunities for post-docs to become entrepreneurs, I'm not convinced that it's a great idea to rely on tying technologically innovative idea exploitation with business interests and skills. I think a really well functioning innovation system is going to enable people to express their innovative ideas, perhaps lucratively so, without having to become a business person, which some people are just not interested in doing or other people are really interested in doing, but are not very good at.

How could an innovation system unbundle innovation from business skills? Well, we have basically some kind of technology transfer process required, and Scott kind of talked about this yesterday. Part of our answer is the patent system, and the patent system is potentially very valuable for doing that, it's one of our key mechanisms for doing that.

On the other hand, as you're all aware, the patent system has gone awry, I think most people would agree, in recent years. So, cleaning up the patent

system, making it work more the way it should work, seems to me a priority in that area.

The other thing that helps, I believe, is if you have a bright idea for a better product, it would be useful to be able to approach a number of different firms who were well positioned to exploit that, and not to be facing a monopsony technology purchaser, because of the dominant market position in the product markets where that innovation might be exploited.

So, those are two things that come to my mind in thinking about how to facilitate that kind of innovation.

DR. LITAN: Joe, just one clarification, because I think both Tom and I would agree, and I think you would agree after we say this, is that our purpose is not necessarily to turn these people into entrepreneurs. It's to pair them with other people because we do not expect them, most of them, to become entrepreneurs. So, we fully agree with you.

DR. FARRELL: Then I agree with you, too.

DR. LITAN: In fact, they are not suited to become entrepreneurs, but the thing is, they haven't thought, in most cases, that their idea is commercially valuable, and that's the key thing, it's opportunity recognition.

DR. STERN: Let's turn it over to questions.

UNIDENTIFIED MALE SPEAKER: Your proposed reform for the monopoly universities, and as I understood it, you said that the home university would still get as big a fraction as it was ever going to get, but if that was true, wouldn't they want to open it up? I mean, it seems like if they don't want to open it, either they're pretty dumb or they're worried that once it gets out of their hands they won't actually get to recapture that share that they think they're supposed to get. Is that a valid concern for them to have?

MR. LITAN: Your question reveals the fact that it is clearly in the universities' interest to adopt the proposal that I talked about, and it's not only my idea, I worked on it with Lisa Mitchell at the foundation. It's clearly in their interest. The question is why don't they do it? And I guess the best answer I've got is I think that for most university presidents, their A-number-one priority is not technology transfer, it's running a university. The TLO is viewed as a profit center, hopefully, and just as university presidents can become captured by their IT department, because the university president knows nothing about IT, I think too many university presidents are captured by their TLO officers, who have a vested interest in keeping their

job, and growing their empire. As a result, too many university presidents are unaware that they are suboptimally maximizing the position of their university. I think a little bit of education can help them.

UNIDENTIFIED MALE SPEAKER: Like every one in the whole country?

DR. LITAN: Yes. Yes.

UNIDENTIFIED MALE SPEAKER: So, there is no valid concern at all that once this --

DR. LITAN: Well, actually, let me tell you something, I can tell you one university, I will be happy to name him, who, if he were here on this panel, would agree 100 percent and knows this to his bones, it's Michael Crow at Arizona State. He knows this, he used to be the TLO officer at Columbia University, he knows the system inside and out. There are going to have to be a lot more Michael Crows to help change this world.

DR. STERN: Joe, did you want to respond?

DR. FARRELL: Yes, I actually had a question for Bob on his complaint about university TLOs. I'm not sure whether I'm caviling about your use of the word monopoly or about the substantive claim. The employees of most firms, if they come up with an invention, have

to get it exploited through their firm. Are you saying that this is always a problem, or what is it that you're saying is different about universities? Is it just that they're less focused on efficiently exploiting their employees' inventions?

DR. LITAN: They're much less efficiently focused. Remember the university, if you view the 80/20 rule, let's look at the 20 who are involved in this. They are basically a research factory, which almost by definition distinguishes them from virtually all private sector firms. As a research factory, they have lots of people. Their TLO office almost by definition is very limited in terms of resources. So, they can only pay attention to maybe one or two technologies or five technologies that they think are going to be the next Google. Meanwhile, there are lots of professors in the queue wanting attention. I know all this because we've gotten many emails from around the country of qualified professors who are complaining that they are ignored by their TLO offices.

I think they are inefficiently hamstrung. Also, to be honest, it maximizes the position and importance of the TLO office to have to be a monopoly gatekeeper. By the way, I think there are some good TLO offices. I can tell you some of the best are MIT's, Stanford's, and

Wisconsin's. In the world that I envision, some of those TLO offices would wind up becoming agents for other universities, and they would take a fee, because they're good at it. I think we have a few more TLO offices in my world, because we have good competition, the good ones would survive. We would probably have some independent agents, too, become involved.

DR. STERN: I'm sort of aware that we are pretty much out of time here, if I'm not correct. So, what I want to do is kind of sum up this, because I think this really, even in a relatively short amount of time, made what to me are some kind of first order fascinating issues. I just want to observe two pieces that we know almost nothing in industrial organization or competition policy about the industrial organization of the research sector of the economy?

We actually even had an entire panel yesterday on innovation in which every single concern that was raised on this panel didn't come up at all. So, there's a real disjunction, I think, between the industrial organization studies of innovation, and this kind of emerging body of research and policy research around the economics of innovation and entrepreneurship. I'll end on a related note, which is to say that I think the kind of pushing forward on exactly how to unleash and figure

out how to take these research factories and try to use them as levers to promote competition in the economy could be a much more center stage, at least in my view, in both the activities of competition policy agencies, which have done, I think, a pretty good job on this, but in particular, on the part of our intellectual property agencies, which really have had difficulty reforming themselves to really use one of the few levers we have to actually make this process work.

I want to thank our three panelists for what I think was a very fascinating discussion and then turn it over to Chris for our closing address.

(Applause.)

DR. ADAMS: So, I think we're going to go directly into the next speaker. Firstly I want to thank Scott for doing that. I think that was really interesting and I'm very glad that we could have Tom, Bob and Joe talk. I think I was very interested in what they thought was going on.

Our last speaker of the day is Howard Shelanski. Howard is my boss, he said yesterday morning that he wasn't my boss, and then yesterday afternoon, he suggested I work on something. He's actually a law professor at University of Berkeley, he also has a Ph.D. in economics from Berkeley as well. I notice he's a

graduate from Haverford College, which is important to me, because they have the only varsity cricket team in the United States.

UNIDENTIFIED MALE SPEAKER: They're number one every year.

DR. ADAMS: And he's certainly the only person in this room who has clerked for a U. S. Supreme Court Justice, Justice Scalia. So, he brings a lot to the table, and I think he's been great for us at the FTC. So, welcome.

DR. SHELANSKI: Thanks, Chris, you brought yourself a two-hour extension with that nice introduction. Thanks, all, for sticking around. We're coming up hard against the lunch hour, and so I will try to make this fairly concise.

When I was approached a week or so ago about giving what I'll call an endnote address, I wasn't sure what I was going to talk about. I decided after thinking about a few topics, to talk about an issue that I think is coming up fairly often, and about which there's a lot of confusion. And it has to do with the whole complex of refusal to deal questions, interoperability questions, denial of access questions, and an interesting phenomenon that's happening in the courts.

So, let me just back up and give a little bit of a preview of the issue before I track through in very reduced form a couple of arguments I would like to make.

I think when one looks at the question of mandated dealing, or refusal to deal by a monopolist and claims by an antitrust plaintiff that that monopolist should have to deal. This dealing can take many forms: The providing of a key input to a downstream competitor, provision of access to a horizontal competitor, all of the issues surrounding interoperability and Microsoft and some big investigation that we also have now before the Commission, and lots of cases.

One hears many arguments about why refusals to deal are bad and harmful. We've seen extremely sophisticated theoretical work on the order of Segal and Whinston telling us about what the net effects of exclusion are by an innovative monopolist, showing us that there may be short-term benefits to that exclusion, but longer term net harms. We certainly can envision the static harms to price and output when somebody is excluded from a key input.

But then once we get to the world of policy without actually mandating dealing, the theoretical, I think, consensus is that there are harms from exclusion by a monopolist, starts to break down when we talk about

policy implementation. Can we find a reasonable way to actually discriminate between those cases of refusal to deal where the exclusion is beneficial, and those where it's harmful?

This gets very difficult. What's happened in the courts is the courts have increasingly adopted a presumption that we're going to have too many false positives if antitrust gets into the business of mandated dealing.

So, what we have seen is a steady movement of the courts from essential facilities cases of the early sixties and seventies through to the Trinko case in 2004. So, it's really narrowing the factual circumstances in which courts in the U.S. will even recognize a refusal to deal claim by a plaintiff.

This trend has been particularly pronounced when it comes to intellectual property. So, the first thing I would want to say is when we think about refusals to deal, we ought to have some reasons why we care. Why shouldn't we just say there is no antitrust claim of a refusal to deal, forget about it, this doesn't constitute the kind of thing that the FTC or the DOJ are going to look at and private courts should stay out of it.

Well, I think there are policy reasons to care

about refusals to deal. Innovation incentives, in fact, might be harmed by mandated dealing, but also could be enhanced by mandated dealing under some circumstances.

There are static harms. There are foreclosure considerations that have come up in some of the papers and some of the discussions we've had over the past day and a half. The trade-offs are uncertain. Whenever we mandate a dealing between deterrents of innovation incentives and improvement of static harms, foreclosure and longer run innovation incentives, but at least in theory we know refusals to deal can harm both. So, we have policy reasons to care about refusals to deal.

Why am I spending any time thinking about them? Well, whether we have policy concerns or not, as a doctrinal reason, the law says that sometimes, though, rarely refusals to deal are illegal. So, there is still some scope for plaintiffs to get in and make refusals to deal cases.

However, over the past 15 years or so, there has been an interesting split in the courts, as they have approached refusals to deal. This split deals particularly with refusals to deal in intellectual property. Here is the general nature of the split, it's a little more nuanced in many more cases, but I think it's well illustrated by the Data General case from 1994

in the First Circuit and by the CSU versus Xerox case in the Federal Circuit in 2004.

What the First Circuit says is that refusals to deal on IP should not be, per se, legal. You can't simply say the product I am refusing to deal in is either IP or IP-protected and get a free pass from the Court. Antitrust rule of reason balancing applies to refusals to deal in IP, but we're going to adopt a presumption that IP represents a kind of innovation or investment that needs enhanced incentives. The enhanced incentives are the legal protection of IP.

Why do we adopt that presumption? Well, that's why Congress has IP statutes. The constitution tells them to have a patent law, but they have framed the patent law in such a way to protect certain inventions. Which inventions? Those that were less likely to get done without IP protections. What is IP protection? It's the very broad right to exclude. So, IP-protected innovations are those that are presumptively more likely to require a broad exclusion in order to have been done in the first place.

When we do our rule of reason balancing, says the First Circuit, we ought to recognize that refusal to deal in IP can be illegal, but put a heavier thumb on the pro-competitive justification for the refusal to

deal, for the need for exclusion.

Antitrust applies, but the rule of reason balance has different weights when IP is involved. That's one line of cases.

There's another line of cases, by the way, that says IP shouldn't matter at all. The Ninth Circuit certainly adopted that viewpoint or may have adopted that viewpoint in the Kodak case, but it's unclear and it's unclear what the law still is in that circuit. The Tenth Circuit has a case that didn't seem to care about the status of the product issues IP, but I think the First Circuit articulates best the view of thoughtful circuits that once you make a nod towards IP you still want antitrust to apply.

The other trajectory in the courts is represented most strongly by the Federal Circuit in the CSU versus Xerox case. And what the Federal Circuit there said was the status of the good in which the monopolist refuses to deal, the legal IP status matters, and in fact it matters an awful lot, because a valid business justification for refusal to deal is that the product is IP protected. The Federal Circuit said that that would be a conclusively valid business justification, that is to say, a pro-competitive justification that trumps any articulable

anticompetitive harm, unless it's not really IP, it was obtained from fraud on the Patent Office, or it's not really an unconditional refusal to deal, it's part of a tie-in claim.

You put those two things together, an unconditional refusal to deal in IP, even in something that good faith could be viewed as IP, even if it would fail an invalidity test in court, that refusal to deal is per se legal.

So, the question I think that comes up for economists and for antitrust policy makers is whether refusals to deal on IP should be exempt from the antitrust test for liability for refusal to deal in other property. Is the First Circuit right that our rule of reason balance can apply to IP, or is the Federal Circuit right, that we should simply exempt the intellectual property?

Or, and then a second question is, should antitrust scrutinize any refusals to deal, or should they be, per se, illegal, regardless of the property at issue? These are two very big questions, but I would like to give two fairly short answers, and I have a lengthy paper that deals with them in greater detail, if you're interested.

So, let me start by saying, I think there is no

reason to exempt refusals to deal in intellectual property from antitrust. I think the Federal Circuit is dead wrong. Just starting off thinking like a lawyer for a minute, there's no legal basis in the IP statutes. There's nothing in the intellectual property law statutes that says antitrust needs to get out of the way, and the Supreme Court has said, we're not going to imply those kinds of exemptions or preemptions too easily. And I would also add that IP policy considerations do not justify an exemption either.

Well, it's not in the IP statute, but we want people to patent things and copyright things and we want to give them strong incentives to do that, and a blanket exemption from mandated dealing under the antitrust laws is a necessary exemption. There's no IP policy consideration in effect.

If I had more time, I would explain why I think when you get into all of the invalidity, and I think Joe Farrell mentioned in the last panel that our IP system has gone awry, the lack of precision and the breadth of our IP system and the 50 percent invalidation rate of contested patents says something, I think, fairly compelling about why IP policy considerations would not justify an exemption.

But as a matter of economics, I think there's no

empirical or theoretical basis in economics, at least not a satisfactory one, for exempting IP. Theory does not provide a good reason to treat IP in a systematically different way from other property. Many kinds of investments require very strong exclusions, very strong kinds of incentives, have very high hurdle rates, that may not involve IP at all. There's very good work that shows why certain innovations, if not IP protected, might not be undertaken, but there's nothing that says that systematically we should treat IP in a different way.

And I think empirical evidence casts substantial doubt on the link between IP protection and innovation. There are the classic Levinthal studies from the mid-eighties that have been updated by Wesley Cohen and others through the nineties and maybe even more recently, that cast doubt on this. The work of Bronwyn Hall, I think, is very compelling in showing that the link between IP protections and innovation is fairly feeble. We get a lot of innovation outside of the IP world. We therefore don't need to give a special exemption for IP to incent innovation.

So, I would say there is no reason economically to allow greater scope for exclusionary practices for IP than for other kinds of investment, at least not in the

systematic way that would be achieved from exemption from Section 2 liability refusals to deal.

So, let me turn to a second question, this is really a big question, and certainly good folks at the DOJ have thought a lot about this had a different viewpoint that was articulated in the now withdrawn Section 2 report. But I believe there's a serious debate to be had over whether refusals to deal on the whole just shouldn't be legal, per se, and when whether we find access or interoperability problems, they should be handled through specific regulation, like in Telecom, the '96 Act has a whole set of access provisions.

I think there are some unsatisfactory ideas, effects, if we relegate all interoperability refusal to deals to sector-specific regulation, just because Congress acts awfully slowly. They will probably get the access terms wrong, and I think problems can arise in many different places and not just where one might choose to litigate.

I don't think refusals to deal should be illegal, per se. Liability for refusals still should be hard to come by, but not impossible. We know from experience that refusals to deal can be costly for consumers in both the short and long term, for both prices and innovation.

Refer again to the Segal and Whinston paper, it's not the only one, there have been a lot of theoretical explorations, both before and since that paper. Just as a nice sort of empirical example of a refusal to deal case that led to some, I think, good consumer price effects, I would point to the AT&T divestiture. Not an uncontested story, but certainly when it comes to long distance prices, I think a fairly compelling one.

So, why do lawyers and courts tend to tilt towards this great fear of addressing refusals to deal? Well, the Supreme Court in Trinko and the other cases and the lower courts have disdained refusal to deal liability on grounds the courts will make mistakes and to deter beneficial innovation by mandating giving them access.

I think this is a valid concern. It certainly has been manifested in extremely slim possibilities that are afforded to plaintiffs seeking refusal to deal litigation. I don't think it should be taken too far, and too far in my view is per se legality for refusals to deal. There are costs from under-deterrence, as have been well recognized within the economics literature and recognized by even the Antitrust Modernization Commission in its report. And I want to also add the

actual experience from courts does not support the broad sphere of letting plaintiffs bring refusals to deal claims.

Glenn Robinson did a nice study looking at all cases from the eighties and nineties where refusal to deal claims were made. Most of them don't get past a motion to dismiss, yet fewer of them get past the summary judgments, and then they're gone, they're out of court.

Now, I am not saying that that's a sufficient statistic. Of course, there are costs to fighting a summary judgment or dismissal motion and there are probably a whole bunch of cases that settle under the radar screen that are costly to the firms that get hit up. But there is an assumption that scholars often make, which is that those invisible settlements are pure social cost; I disagree. A lot of those settlements involve cross-licensing agreements that bring benefits to consumers and free up opportunities to innovate.

So, I think there's no rationale for assuming that the hidden settlements obscure social cost. They may well be beneficial and I actually give them just the neutral weight.

So, I think then looking at the observable cases is not a bad way to get a sense of whether the courts

have run amuck and are letting a lot of bad claims go through. The evidence is very slim. I've updated Robinson's data set for the cases from 2000 through 2009, and it's really a very small handful of refusal to deal cases that, again, have even made it to the trial stage.

So, what we don't see plaintiffs extorting huge settlements and huge costs and, in fact, when the Antitrust Modernization Commission asked for evidence of deterrence or high costs from these kinds of cases, nothing was submitted. I think that was a pretty good opportunity for business to come to the table with a lot of evidence that they were being thwarted by these bad monopolization claims and it was never produced.

So, let me just conclude with some open questions and suggest how economists can help with further policy and thinking on these topics. How should we define or identify harmful refusals to deal? I still don't think we have a very well-articulated theory in the case law. I would love to see some more thought on this. Empirically, how prevalent are harmful refusals to deal? I don't know how we get at this, but I think some clever empirical thinkers could come up with some metrics. I certainly think we could look at the next suggestion, which is whether or not retrospective

analysis can tell us whether cases that have been heard in the past where refusals to deal liability have been opposed have helped or hindered competition in innovation.

If these cases had been a waste, well maybe that's in favor of per se legality for refusals to deal. If we have seen entrenchment and serious costs from not intervening, it weighs in the other direction.

Finally, does IP warrant a stronger presumption of innovation deterrence for mandated dealing, even if not an exemption of refusals to deal on our IP from Section 2? This goes to the First Circuit's rationale on Data General. Are they right to put a heavier thumb on the scales of pro-competitive justification for exclusion and refusal to deal on IP? Or should there actually be no special status at all or special weight for IP at all in refusal to deal cases?

Just a handful of questions. I think it's worth thinking about. This is going to be an ongoing policy issue; we're going to see continued cases here at the Commission of what amount to refusal to deal claims and interoperability claims. We need to know a little more about how to handle them and I think the courts have made something of a hash of it. So, it's going to come back to the Supreme Court, eventually, because of the

circuit's split on the treatment of IP, it would be nice if we could give them more guidance on how to rule.

Thanks.

(Applause.)

DR. ADAMS: Let's just take maybe one or two questions. Do you have a question?

UNIDENTIFIED MALE SPEAKER:

Howard, if the potential liability for refusals to deal in IP were broadened by refusals to license patents, for example, to what extent are you concerned or how would you manage a marginal incentive to substitute trade secrets for patents? And, of course, that substitution would mean the information made public in a patent filing would not be made public, and that could at least have some effect on diffusion of new processes.

DR. SHELANSKI: That's a great question. Obviously, as you change the legal regimes as they bear on one form of protection, you may get substitution on the margin, which could be quite a big margin for other kinds of things, like trade secrets. There are some drawbacks to trade secrets because you don't have disclosure, but there's a big benefit to trade secret, which is the ability to reverse engineer and completely replicate the technology.

So, I think there's a fair trade-off there. So, I don't worry about it so much. Trade secrets tend to, although not always, have greater fragility. Eventually, firms are pretty confident that they're going to get reverse engineered.

Some very high-tech firms use trade secrets now because it's not worth going through the patent process for a technology that's going to be obsolete in two years.

When you talk to Hewlett Packard, they have tons of unpatented, seriously valuable IP, but they never patent it. They just protect it by trade secrets because that will get them the 24 months they need of using this before it obsolesces, and then it gets reverse engineered and it's usually reverse engineered to the end of that life cycle.

So, I don't worry about the marginal substitution because I think in a lot of cases where the firms would move to trade secrets not to disclose the IP, it has that compensating benefit of reverse engineerability, if that's a term.

I almost worry more about the other side of it, which is we already have too much junk being patented. The Patent Office is overwhelmed. If you start to grant increased exclusionary privileges, just because of the

legal status of something as IP, you're going to have this marginal incentive to drive things towards patenting that perhaps shouldn't have been, with possible harms and certainly some costs to the system.

Dan?

DR. O'BRIEN: So, when you talk about refusals to deal, if you say, well, you can't refuse to deal, you must deal, it sort of inevitably morphs into a question about dealing at what price, and every time I think about this question, it always comes for me to, well, we're basically saying we should be using the antitrust laws to, in certain instances, to effectively regulate, price regulate, industries that we think are not adequately regulated by competition.

I'm wondering about whatever mechanism, the market mechanism, or maybe by regulators. To what extent is that an accurate description of what you're thinking?

DR. SHELANSKI: So, there are two questions there. One is, is this regulatory problem, the court as regulator, the agency as regulator, a prohibitive problem? I think it's a big problem, but not prohibitive. Often there are benchmarks that can be used or in some cases zero price could be appropriate. I have another suggestion, which is simply this: In a

lot of those cases, in fact, the most compelling ones, you're probably dealing, not all of them, but in a lot of them, you're dealing with a downstream, vertically related competitor. So, you have a vertically integrated firm refusing to deal with your downstream rival, and that's what they're complaining about.

In that case, my answer is sell at any price you want, just be aware that if you get hit with a predatory pricing claim in the downstream market, the attribution will be the price you're charging to the outside downstream competitor. The cost attribution in a price cost estimate for your downstream product.

We don't have to worry about pricing then, and it's at least a second best, but a possible solution for dealing with the problem. It requires recognizing margin squeeze in a particular way. And I know that makes you lose your appetite for lunch, but there's something that I think there are solutions there.

Your broader descriptive claim, or question, isn't it really the case that these are just industries where we think broader regulation is needed and we're unhappy with monopoly more broadly. I'm not sure that that's the case.

If that's the case, and that's an empirical question, then really we ought to just move to

sector-specific regulation and abandon refusal to deal as an antitrust claim. I don't think that's the case. I think it happens fairly ad hoc and lots of industries that wouldn't want to have broader regulation, but we might need a licensing remedy or we might need some kind of mandated dealing.

It's rare. I think we need to specify the conditions a little better than the facts at Aspen, because that's simply ad hoc and Aspen is a preference, but I think it's still an open question.

My instinct is to say, yes, there are facts we can identify that would be better than not to mandate dealing. And moreover, it's a manageable problem where we do have to articulate the terms of the deal. But I go back and forth on this, and that's why I phrased them as open questions.

Do we have time for one more or are we done?

(No response.)

DR. SHELANSKI: I'll talk to Pat afterwards.

DR. ADAMS: Definitely we're done.

DR. SHELANSKI: Other than Dan, who has lost his appetite thinking about refusals to deal, we're probably all hungry, so thanks.

(Applause.)

DR. ADAMS: So, I would just finish up, I want

to thank Paul Rothstein for co-organizing this conference with me. I'm very grateful for all his help and ideas. I think we were very lucky, again, to have a fantastic scientific committee. Kyle just left, we had Aviv and Marianne who helped us a lot, and Scott as well. It is fantastic to have those people helping us. Most importantly, I wanted to thank Laura Kmitch. I think she's outside somewhere; she's just done a huge amount of work putting this thing together and make it all run very smoothly. So, I'm very grateful to her.

Thank you all for coming, and enjoy your lunch.

(Applause.)

(Whereupon, at 12:39 p.m., the workshop 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 20, 2009

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: 12/4/09

SALLY JO QUADE

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.

SARA J. VANCE