FEDERAL TRADE COMMISSION

ECONOMICS OF INTERNET AUCTIONS CONFERENCE

Thursday, October 27, 2005

9:00 a.m.

Federal Trade Commission

FTC Conference Center

601 New Jersey Avenue, N.W.

Washington, D.C.

FEDERAL TRADE COMMISSION

Contents

P R O C E E D I N G S

\- - - - -

INTRODUCTION

DIRECTOR SALINGER: I'm Michael Salinger. I'm the Director of the Bureau of Economics here at the FTC. And I'd just like to welcome everyone and thank you for coming to what promises to be a most interesting conference that will touch on three of the four major missions of the Bureau of Economics.

In our consumer protection mission, fraud and internet auctions has emerged as a major source of consumer complaints. And this morning we'll discuss that topic. And then this afternoon we'll talk about competition in internet auctions, a topic of potential importance in our antitrust mission.

And the Bureau of Economics also has as part of its mission to conduct academic research. And, so, we'll finish up today talking about the use of data from internet auctions in our research subjects. So, we'll end up talking about research and how we can use data to do better research, to inform our mission to understand consumer protection, antitrust and regulatory intervention issues.

We're also really fortunate today to have -- and honored to have Hal Varian to give our keynote address at lunch. So, we're very pleased about that.

I have not been here very long. And this conference was well in the works before I got here, so I think I can say the following without being self-serving. To put on a conference like this of this quality, to cover such a range of practical issues and issues of academic research, is something that not many organizations can pull off, and I think the ability to do it is what makes the Bureau of Economics at the FTC such an amazing institution.

I'd like to thank Chris Adams. This conference is his baby.

(Applause.)

DIRECTOR SALINGER: He conceived of it and he did much of the organizing of it. I'd like to thank Denis Breen, who saw the wisdom of putting the conference on and who shepherded it through the system. We'd like to thank the Chairman's office for giving us the go-ahead on this; and our colleagues in other parts of the Commission, most notably the Bureau of Consumer Protection for your help and support in putting the conference together.

And these conferences administratively are a real challenge to put on, and I'd like to thank Maria, who just disappeared, and all the excellent administrative staff for doing such a fine job. Having stayed up too late last night to watch the end of the ball game, I am particularly grateful for the hot and good coffee that we had this morning.

Having mentioned the baseball game, I will -- from a long-suffering -- former long-suffering Red Sox fan, congratulate the long-suffering White Sox fans in the audience.

So, with that, I think we can have a great conference, have some fun. And I'd like to introduce our first speaker, Professor Patrick Bajari from the University of Michigan.

PRESENTATION: INTRODUCTION TO ECONOMICS OF INTERNET AUCTIONS

BY PROFESSOR PATRICK BAJARI

PROF. BAJARI: Okay, so what Chris asked me to do this morning was -- since I didn't have any new research on eBay that was ready at this instance, he suggested that I come up with an overview of some work people have done in the field and some open questions that researchers might care about. And he asked me to try and make it accessible to a general audience so I don't scare all the industry people in the room.

So, what I decided to do was basically talk about just some of the empirical regularities that the, you know, by now 50 or 60 empirical papers about eBay have found. I think it's worth noting that we actually have empirical regularities in this literature. If you think about lots of literatures and empirical industrial organization and you ask yourself what are the regularities coming in through the literature, are the regularities from differentiate product demand estimation?

It's a little hard to come up with an answer on what that might be. One of the nice things about eBay is, you know, while we may differ on our interpretation of some things we find, there are a certain set of facts that kind of consistently come out in research. These facts are novel; they're somewhat robust to the peculiars of our econometric methods. And I think they've changed the way some people thought about auctions. So, you know, I think there are open questions, but there is some degree of success in this empirical agenda.

I think one of the reasons we've had this success, or some degree of success, is because the data's so good that when we go look at data on eBay or for that matter in other eCommerce sites, we see what the consumers see, which is a lot better than many of our data sets. You know, if I think about the data and differentiated product demands, there exists a fairly large window between what we see and what the consumers are actually doing. So, I think we can, to some extent, thank -- thank the quality of the data for this.

So, I'm going to go through a few of these regularities that we seem to find and then suggest some questions where I think we've got some -- where academic researchers have some work to do. And these questions may seem like pointed academic types of things, but I do think that getting these questions right is important to the types of policy issues that you were bringing up about fraud and so forth.

So, when I was in graduate school, the model of bidding I learned was Milgrom and Weber. And we learned that, you know, auctions are these beautiful mechanisms that allowed markets to function in the presence of private information, but under fairly robust conditions, you know, with different auction designs, that the person who valued the good the most would win it; and that their market power would be limited; that the margins would go to zero as things became reasonably efficient.

And there's this sort of extremely beautiful theory that says, you know, what people should go out and do in an auction is keep bidding, as long as their valuation exceeds the bid. So, that's why you have private values, which for the non-academics means you aren't worried about fraud and getting burned, you sort of know what you're purchasing; or common values, which is a setting where we're more concerned about adverse selection, uncertain about what we're getting. You should just keep in an auction until your valuation falls below the standing bid.

Well, that was not such a great way to describe what went on in eBay. These sort of -- when you look at the data and look at these standard models, there is really not a lot of comparison between the two. And empirical researchers pointed this out. We've helped to spawn a new little theoretical literature as a result that has fought through.

Certainly we've been through the dynamics of auctions and trying to come up with richer and more relevant models. So, you know, one of the things we saw was that things like sniping behavior, you know, we saw this flurry of bids that would occur at the end of the auction, particularly from bidders that were more experienced. And people found this over and over and over again. And, you know, the simple fact of pointing out that, you know, people tend to bid 2.5 times in the auction with particularly density at the end, and the people bidding at the end tend to win, I think was a challenge for some of those theoretical models.

And, you know, we look at that and the theorists have to start scratching their head and asking, you know, do we have allocated efficiencies here? Does this look like collusion, you know? What in the world is going on? We've all differed about our interpretation of these things, but I think pointing it out again and again in a robust way has stimulated economic thinking, so people came up with the very serious issue, how a

-- you know, fairly benign looking ascending auction could support collusion. You know, people thought about behavioral theories of this and so forth.

Another thing we seem to have found are some regularities about reserve prices and entry behavior, you know, the way, when I was in graduate school, I thought of reserve prices as, you know, this fixed number of bidders, maybe there are enough bidders for competition, therefore, the seller could use the reserve prices and mechanisms to induce some extra competition in the auction.

That was not a particularly good theory of what was happening in these markets, but people would say over and over again in the empirical literature that reserve prices seem to be inherently linked to participation, they are linked to your likelihood of selling the object, and that bidders are -- sellers are sort of trading off, getting more people in the auction versus actually being able to complete the transaction for their choice of a good price. And I think a lot of people found that secret reserves tend to discourage participation.

You know, tell me one theory that existed five years ago that thought of any of these things. They aren't out there. I looked hard. Maybe there were -- I overlooked them, but I don't think they're there. So we've found these things fairly robustly, and I think they've changed the way we look at these mechanisms.

And this is interesting to us as economists because we care about market design, about how to set up markets, so they function efficiently and properly, how to advise sellers on how to maximize their revenue and design auctions correctly. And, you know, we've really had our differences in terms of our interpretation of why these things are going on or some of the -- some of the effects of reserve prices, but because of the high quality nature of this data, I think we found this again and again in quite a robust way.

The third thing that I think we've found are some indications that adverse selection and information asymmetries are important. And this is useful, you know? I came from a big macro school when I was a graduate student, you know, like this Minnesota/Chicago type tradition, and I talked about information asymmetries, I think that's just second order stuff, that it doesn't really play much of a role in markets. And, you know, if you went and really scoured the evidence, the set of cases where we had, you know, really good documentation on, you know, here are information asymmetries, here's where they matter.

It was a little thing. So, there have been at least, you know, four or five different papers that have pointed this out. So, Ginger Jin's worked with Andrew Kato on baseball cards, showing that, you know, it looks like there's a possibility of these markets not working correctly; Pai-Ling's work on sort of showing that when bidders are more uncertain about what they're getting, they're paying less, which was -- you know, this is -- a lot of you have probably seen this draft before, but this is -- when I go and teach graduate industrial organization and I talk about, you know, evidence for the winner's curse, this is probably as convincing a piece of evidence as I've seen. This has on this axis -- we have from human subjects, how -- a measure of how uncertain they are about what they're getting; and here's the price. And this sort of says, you know, you go from the 25th to the 75th percentile of how uncertain people are, the price just plummets, 60, 70 percent. That is a much more convincing piece of evidence than regressing bids on the number of bidders like we've done in other places.

And then, finally, you know, a lot of people -- probably the most studied question is about reputation. We have all these elegant theories of reputation, talking about how reputation is this beautiful commitment mechanism that can help markets function in the presence of adverse selection and moral hazard, that I won't screw over my fellow man because there's this option value of having a reputation for being a decent and noble human being.

But, you know, once again, the theoretical evidence about how important is reputation, where is it in the markets, where is it in the data, was much thinner. So, here's where the empirical people have come in. I've seen I think -- Ali, he did the hard work, I didn't. He tabulated 15 of these different studies and then we typically found that the signs pointed in the right direction, that people who did good things got rewarded for it in the marketplace, people who did bad things got lower prices. This is probably about as good of evidence as we've seen about the importance of -- kind of quantifying the importance of reputation in market transactions.

So, in the last four minutes I've got, let me tell you three questions that I think are interesting. The first one is, we don't know how to look at these bids yet. We don't know how to understand the richness of the dynamics that are in the data. And this may seem like a pinheaded academic question, but unless we understand where the bids are coming from, we cannot answer fundamental questions like, is the market operating efficiently, is the person who values the object the most winning it?

So, let me give you just a cute little example I found yesterday on eBay. This is a set of bids I found for iTunes. So, it was 50 iTunes. The seller would go and email you the codes for the iTunes after you won it. As an economist, I found this set of bids for 50 iTunes to be really confusing. So, you know, iTunes have a pretty well-known market price, they're 99 cents each. So, 50 times that should be like $49.50. First of all, the top line here should be disturbing for anybody who cares about market efficiency. The winning price was $40. So, $49.50 was only worth $40 in the marketplace, which, boy, that -- and this is not isolated either. You see this more than once.

And then what are people doing here? I mean, they start out bidding $5 and they're ratcheting it up and only bidding seriously at the end. I have a hard time saying this is differences in private values. I mean, if I'm going to use all these things, it should be worth $49.50 to me. When I go to cheaper iTunes, I find the same thing. The bidding is still a bit mysterious, and until we know how to say something intelligent about examples like this, I think our ability to diagnose efficiency and ask just basic economic questions here is limited.

A second open question is -- and I think this relates to a lot of the regulation types of questions that I've heard people at the Federal Trade Commission talk about -- is the role of asymmetric information. We have -- when you talk about things like, are people getting ripped off, is consumer fraud taking place, well, you know, in any used good market, we're going to have some problems. People are going to be upset about used goods. But how bad is that? Is this functioning in a sensible way? Are the prices reflecting people's uncertainty about what they're getting? You know, basic things about the theory of asymmetric information, when markets are working right, you know. In Milgrom and Weber's set-up, people might be upset about what they're buying, but, you know, the market has certain nice properties.

We don't really know the answers to those questions. I think if we're going to get to the answers to these questions, we need to do two things. First of all, we've got to think about the identification results and the auction literature, seriously. There's a paper by Quan Voung and Jean-Jacques Lefonte in the AER that said, you know, you can't really tell about the difference between common and private values in the data. For those of you who are industry types, this means if the FTC goes and says, well, there's all this -- you know, people are getting ripped off, this result of Voung's says, how are you learning about that? He has a constructive example where that's problematic. We can't be doing things like saying, the market is functioning poorly, without talking correctly about identification.

I think the other way we're going to get at this is by merging field and survey data. The only way we're going -- our problem in auctions is we don't see people's private information. But we've got these guys' email addresses. We can go talk to people who are bidding in these markets and we need to incorporate that type of information to get a grasp on the omitted part of our models.

And then one last point, and I'll finish up in 30 seconds here, is we still need to think about basic questions and demand estimation. A lot of empirical work in eBay is regressing the winning bid on the characteristics. There's this old econometric paper 20 years ago that tells you why that is an upwardly biased estimate. The intuition is simple. You're looking at what the person who won the object was willing to pay at the margin. That is not the valuation of all the other people. This is an upwardly biased estimate. We are not doing our demand estimation correctly here.

So, when we do things like talk about, well, here's the value of a reputation from regressing bids on seller characteristics, no, it's not. That's an upwardly biased estimate of it. Now, it's all fine and dandy for me to be smarmy and say that, but it's a hard question to know how to do this right. This market has a lot of interesting and hard features to think about. We don't know exactly what those bids are. There are all these minimum bids and reserve prices that complicate our ability to study this.

And the stuff about demand estimation, you know, when regulators or academics go look at the market and say, oh, it's not functioning correctly for this or that reason, you know, ultimately, their answers depend upon doing this step right and we don't quite know how to do this yet.

So, I'm going to stop there, and thanks a lot, Chris, for inviting me.

(Applause.)

PRESENTATIONS: ECONOMICS OF FRAUD AND INFORMATION

DR. DURBIN: All right, so we'll start with our first set of speakers now. We'll have Pai-Ling Yin, Luis Cabral and Ali Hortacsu, and I guess I'll ask them to come up and take seats here, or maybe for the purposes of PowerPoint, it's better not to do that. You can stay in the audience so we'll be able to see what's going on.

Just a couple of quick announcements, first of all, the AV people have told us that, apparently, wireless devices, cell phones, Blackberrys or whatnot, may interfere with the operation of the microphones. So, at least please be aware of that if you notice the microphones are getting fuzzy or whatnot and you're typing on your Blackberry. That may be the reason.

Also, we'll have for each of the speakers now about 20 minutes for the speaker to talk, five minutes for the discussant and then some time for questions. As far as questions go, because we're trying to transcribe this, we sort of want everybody to be talking into a microphone. So, we'll have somebody with a roving microphone. So, please raise your hand and we'll deliver a microphone to you and then we can do the questions.

So, we'll start with Pai-Ling Yin.

PRESENTATION: INFORMATION DISPERSION AND AUCTION PRICES

BY PROFESSOR PAI-LING YIN

PROF. YIN: Okay, all right, thank you very much. I think I'm going to have to stand here in order to move the slides. But thank you very much for coming and thank you, again, to Chris for organizing this.

I’m going to talk about information dispersion and auction prices. Pat did a great job of establishing two areas where we find that there’s incomplete information in auctions, and in particular, I’m going to be studying the area of eBay.

So, the first type of incomplete information that is most often cited is this information asymmetry problem. So, sellers may know more about the good than bidders, and as a result, sellers may take advantage of this by trying to deceive the bidders.

Now, a second type of incomplete information is actually information dispersion. So, a lot of used goods are sold on eBay and even the new goods are sometime sold by people who have made their own goods. So, it may even be unclear to the seller what the value of the item is. So, you may have these different dispersed private signals about the value of the goods that are dispersed across both buyers and sellers. This may be a result of maybe these bidders and sellers having some differential experiences with that good.

So, when we look at information dispersion, one of the problems is that if bidders are unable to account for this information dispersion then the winner may suffer from the winner’s curse. What that means is that they may win the auction, but at a price that’s higher than the common value of the good, if we’re talking about common value goods.

So, one way to solve this dispersion problem is that we may have -- we allow the seller and the auction to provide an auction description.

UNIDENTIFIED MALE: Can I interrupt? Can you stand back behind the (inaudible) so we can (inaudible)?

PROF. YIN: Oh, sorry, okay. All right, let me just stay here then. I'm sorry, too much training in the HBS classroom that we have to roam around the classrooms. **(Laughter).**

PROF. YIN: So, whether or not these bidders do actually account for the winner's curse, however, in eBay auctions is actually an empirical question because there have actually been mixed results in experimental and commercial studies about the bidder behavior. Now, one way to solve the information asymmetry problem, for instance, on eBay is that they create a feedback mechanism that creates a reputation for the seller.

When you combine the idea of information dispersion and information asymmetry, then you can get out the question of, does reputation affect prices through the bidder's perception of dispersion? A lot of the literature that has looked at the feedback mechanism has focused on a price premium. But, basically, what this work is actually going to focus on is whether we can tease apart the effects of information about the good versus information about the seller. So, in fact, are these complements or substitutes?

Now, the way I'm going to try to test this is, first of all, to look at the theory and figure out what are the implications of common value auctions with Nash equilibrium bidding, so basically rational bidding behavior. That distinguished that model from models of common value auctions with naive bidding behavior where people don’t take into account the number of bidders that are in the auction and don’t take into account the information dispersion, or private value auctions where people are just bidding their own value.

And in order to do this, though, I need a measure of information dispersion that isn’t already going to assume rational bidding behavior and isn’t already going to assume common value auctions. So, I’m going to argue that I successfully construct such a measure by doing a survey of people’s valuations of items, and I’m going to use this survey and the data that I collect from the actual eBay auctions that we all see to test these implications. What I find, basically, is that prices are consistent with this Nash equilibrium behavior in common value auctions.

In addition, I then correct for any measurement bias that might be in the survey, and I use it to create counter-factuals in order to estimate the amount of winner’s curse in these auctions. And what I find, vis-a-vis, the different types of information that could be provided by the seller that, in fact, sellers’ reputations are complements to the information that they provide in an auction. So, basically good reputation sellers have an incentive to provide a lot of information because they get a lot of return from that in the final prices. So, we see a reason why there’s an incentive on eBay for sellers to actually reduce the amount of uncertainty in the entire market, and perhaps, this is the reason why eBay may have efficient trades or promote efficient trades.

So, just to set up, first of all, I’m just going to go over the theory model that’s going to give me this testable implication that’s going to distinguish rational bidding common value auctions from private value auctions or common value auctions with naive bidding.

So, basically, we have a single indivisible item with unknown common value, which I'll denote V and it's going to be sold at a price, P. There are end bidders who are indexed by I and they all know the distribution of the common value. They get a private signal about the common value, X_ , and they all know the distribution of that signal, conditional on the common value.

The assumptions are that we have risk neutral bidders with a utility, the value of the item minus the price of the item. And then I make two other assumptions that -- first of all, that these distributions are continuous and that they have first and second derivatives, and that these distributions can be characterized by a location parameter and a scale parameter. Now, actually, these last two assumptions can be relaxed. Basically, I use them in order to conduct these counter-factuals and actually sort of simulate what theory would predict are the correct prices in these auctions.

So, how do we think about information dispersion then in this model? So, imagine that I know that a Hewlett Packard computer has a very noisy fan. So, relative to the value of the item, an actual HP computer, I may actually have a lower signal of what its value is because I know this noisy fan aspect. Ali, on the other hand, may know that the HP computer is very easy to wire and very easy to install more memory. So, his signal may be farther up here from the true value. And then

Luis, for instance, may not know any of this information, so he might be closer to the common value, somewhere in between us.

Now, as I said before, one way to reduce the information dispersion between us is to actually tell all of us in an auction description that it's an HP computer, it has a noisy fan, but it's very easy to wire and add information. So, what that does is it actually brings all of our signals closer together. So, this is how we want to think about reducing information dispersion.

Now, what I want to do is actually then test the implications of this theoretical model and see if I can observe in the data whether or not the eBay auctions that I'm studying -- and in this case, I'm going to use a sample of computer auctions in eBay -- to see whether this common value model with rational bidding is actually the one which we should think of when modeling these auctions or whether a private value model or a common value model with more naive bidding is appropriate.

And what I'm going to use is a result from Milgrom and Weber which says that if you reduce the amount of -- if you publicly reveal information in a second price auction, then price will go up. I've translated that to fit this sort of model by saying that, well, if you reduce the amount of information dispersion for all bidders, then price will go up.

So, in this next chart, the way one should read this is that -- think about the first column as, if I observe these patterns in the data then it is consistent with one of these models. So, for instance, if I see price is actually decreasing as you increase the amount of information dispersion, increase the amount of uncertainty, then that's consistent with a common value Nash model, and it is not consistent with a naive model over a private value model if you make the assumption that they're -- or if you know that these -- distribution of these information signals is symmetric. So, think about a normal distribution.

Now, if you don't know whether or not they're symmetrically distributed or non-symmetrically distributed -- so, for instance, think of a law of normal distribution -- then you're actually not going to be able to get identification for this. It could go either way.

However, what we also know from the theory is that although prices in the limit will increase with the number of bidders. It is true that for common value auctions, prices may go up or down away from the limit. So, as you have just a few number of bidders, you might see both patterns. So, if we were to actually see prices go down with the number of bidders, this can only be consistent with common value Nash, because if they were private values, as soon as you add more bidders, you're always going to have the price go up with the number of bidders. So, that's what comes out of the theory.

The other aspect that I want to add is just thinking about, well, how do we include information asymmetry in this reputation idea in here? So, for instance, let's let R denote reputation or we can think of it as the credibility of the information. So, a bidder may see the information that the seller provides, but how does that bidder know whether to believe that information? So, that's what's going to be captured by this feedback mechanism. And then what we want to think of, this Sigma X conditional on V, is that actually that's the reputation-free dispersion of information in the auction or the reputation-free uncertainty that the bidder views. So, this is the information about the object, that's the information about the seller.

So, R can enter the price in two ways. First of all, R might actually -- the reputation of the seller might actually shift how much you think the average value of the item is. However, when we think about -- so, this is R shifting.

(Laughter.)

PROF. YIN: And then, if we think about reputation affecting the dispersion of information, then a reputation might change that dispersion from looking like that to looking more like that. So, if you recall the example of bringing these signals closer together. If they're credible, then you actually believe that. If they're not credible information, you're not going to be able to bring those signals together.

So, Akerlof and Milgrom and Weber, themselves, also talked about the importance of reputation as one way to prevent sellers from trying to deceive the bidders. So, you know, the basic thing that we might be able to test in this data, as many people have done, is does price increase as you get a better reputation?

However, if we think about, you know, the interaction effect of these two things, then we can also test this idea, that perhaps prices are decreasing as you increase the amount of dispersion at an increasing rate if you have a better reputation. So, the way to think about it is this. If I have a good reputation and I provide a lot of information, I should get a lot of return on the price that I get from the bidders.

However, if I have a bad reputation and I give it a lot of information, then probably the bidders aren't going to believe what I have to say, right?

Now, the very interesting part of this result is it also would imply that if I have a good reputation, but I provide very little information, so I leave uncertainty out there, then probably I'll actually get a negative hit to the price, and the reason is because the bidder is looking at me and saying, well, this person has a good reputation so I'll believe the information, but there's some reason why they're being very vague about the type of stuff they're giving. Maybe they're trying to hide something because of this to protect the reputation effect. So, we're going to just see if this occurs in the data.

So, my sample of eBay auctions are 222 auctions, collected on two different days as opposed to this space -- June 24th, 2002, and July 12th, 2002. They were in the PC Desktop category, recent Pentiums, and so, this just gives you the average, the median and some range values. So, you can see there's actually quite a bit of range in both the prices, as well as the feedback score of the sellers. And I've broken it down, also, into how many of those scores were negative scores.

Now, my survey is actually a survey of all of those conducted on all of those auctions. So, what I did was I took the auctions, I stripped out any seller information and just left the auction description, and I put a survey out on the web and had friends and their friends answer this survey with the opportunity to win a prize. And basically, the important thing to know is that these people who were surveyed are not at all related to the actual bidders in the auctions. So, they were not the auction participants; they're actually just these random other people.

And what I asked them is if you had a friend who said, I want to buy this computer, what is the most that you think that she should be willing to pay? And I denote the average of these estimates as a capital V and the standard deviation of the values as small SD. So, again, this would be the reputation-free information dispersion.

And so, here are the results of the survey and sort of, you know, one thing I'd like to just point out is that I did get quite a few responses per survey, so that was good. In addition, the value of the survey respondents, in general, was actually much higher than the average value of the prices that we actually observed. So, that's just something to note.

Now, the thing is though -- so, the first question you probably have is, well, how do I know that I should trust this survey? These people aren't related to the auctions at all. So, one thing to just test the validity -- well, first of all, let me give you an example of whether or not there's actually dispersion in these markets. So, if you notice in this example, this is a Gateway computer and the seller describes that the computer appears to be dead and the hardware components may still work, but they're, again, not sure. So, this was the auction that had the highest amount of information dispersion. So, there is this level of dispersion that actually exists.

Now, the way to test whether the survey kind of gets it right is to then say, look, let me look at two auctions from my survey who the survey respondents said had about the same value. So, in this case, I'm -- but let me look at two computers that have the same value according to the survey respondents, but one of them has a high standard deviation, so high information dispersion, and one of them has low information dispersion.

So, this is the high information dispersion auction. Some of the things to note is that it says it's a computer system, but it's not clear whether the system includes a monitor or a keyboard, and then they also have this line down here that the computer is similar in style, but not identical to the unit pictured above. So, who knows whether this is this computer and, you know, unless you know who "Wham" is or have an experience with them, you're not really sure how much this computer might be.

The low information dispersion example I have here, I actually cut out a bunch of the details so it would fit on a slide, but basically it had a picture of the item, and one thing to note is that -- sort of that I cut out here is that most of the basic characteristics were about the same, same hard drive, same speed, same memory size, et cetera. But the interesting detail that this seller provides is that this computer works fine when hooked up only to a monitor, printer and speakers with no other additional hardware options. But when I connect the zip drive and scanner, the computer starts to have problems. So, the seller is being very honest about exactly how bad this computer is.

So, you know, I think this just shows that -- and if you believe that people should be reacting to the information dispersion, what we find is actually that the first computer sold for $55 and this computer sold for $95. So, people are reacting to this more detailed revelation of information. And I think you can understand why, then, I'm also interested in doing these surveys because we capture a lot of these idiosyncratic semantics in the auction description that we don't get from just looking at the eBay data.

So, all I did here was just assume that the survey kind of gets it right in general, so that it's correlated with the actual values of information dispersion, number of bidders, et cetera. So, basically, what I find is the general result holds that -- the general results are consistent with common value Nash equilibrium or common value auctions with rational bidding.

So, what we find is that, in fact, prices are decreasing with the standard deviation, so the amount of information dispersion. And we find that prices are decreasing with the number of bidders. So, even though these potentially are biased measures, as long as they're correlated, we're getting the signs right.

Now, one thing that I do -- since I don't have too much time, I'm going to sort of glide over these. I do a process for correcting for the errors in the survey. So, I basically put in some free parameters, and in my survey, I actually collected information about the background of the survey participants, and some of them I know are experienced with eBay; others I know are not experienced with eBay, since these are random people. And what I do is I model the experienced people as probably much more like the people who actually participate in the auctions than the inexperienced people. So, what I'm going to do is model sort of some parameters to allow for how much difference there is between the experienced and non-experienced people.

So, when I throw that in, basically, what I get is that the -- well, these results don't actually mean anything to you since I can't go through the model in the amount of time. But basically what I get is that, in general, the survey participants, both the inexperienced and the experienced ones, basically get the scale right of how much these auctions are worth. But often they're off by a level. So, for instance, the experienced guys overestimate by about $83 and the -- the experienced guys overestimate by about $27, the non-experienced survey participants overestimate by $83.

I tried various things about instrumenting for the number of bidders, but basically the results are not significant because there aren't very good instruments for that. But they, again, are consistent with those initial OLS regressions that I just showed you.

Now, this is the graph that I'm most excited about. So, basically, what I did here was I took all of the eBay prices in my sample of auctions and I just lined them up, and then I took the survey data, with these slight corrections for the amount of bias that might have been in them, and I basically made the assumption about a law of normal distribution, and I tried various other types of distributions. I just simulated, what would auction theory predict are the prices that people should be bidding if these are the -- if the survey data that I have is correct?

So, what I find is actually there's this just, I think, very impressive match between what the eBay bidders are actually bidding and what are the predicted NASH equilibrium common value prices, based on the dispersion of information -- private information signals that I got from my survey. So, it's true that they're kind of overbidding a little bit, but the fact that the pattern matches, I think, is evidence that, look, bidders know that they should take into account the number of bidders and they kind of adjust for it in the right manner, maybe not at the right magnitude, but they sort of get the idea that as there are more bidders, I need to be more careful. So, therefore, they're taking into account information dispersion.

I did some -- then some counter-factuals about thinking about how do we break out reputation from information dispersion. So, the way to read this is that this is the direct effect from decreasing your dispersion or increasing your score or reputation. This is the interaction effect. So, this is about the credit. How much do I believe that information or -- based on my reputation and then the total price effect.

So, the way to think about this is actually that if you have a -- if you compare the baseline of a seller with no reputation and a medium level of information dispersion, if that seller actually increases their reputation to be the medium level of all of the reputations in my sample, which was like a feedback score of 68, then they'll probably get an extra $6 approximately for the price of their item. However, if that same seller were to reduce the amount of dispersion down to the medium level down to about the lowest quartile of the sample, then they actually could increase the product price by I think it's $46. So, there's quite a big difference in sort of what you get back from decreasing the amount of information dispersion; i.e., providing a lot of detail in your product description, versus what you get from reputation.

So, it's clear that these things are, in fact, complements, the information on the product as well as the information about the seller and the seller's reputation, and that eBay markets account for a significant amount of winner's curse. So, there is just that measurement between how much people are overbidding and the actual prices in the auctions. And I think, again, the exciting thing is that you can use this survey data to provide you with extra information that allows you to test sort of what kind of auction setting are we in, are people behaving rationally, as well as, you know, do some of these measurements support the idea that maybe people are kind of getting it right on eBay?

So, that's it. Hopefully, I'm on time. Thanks.

DR. DURBIN: So, as a discussant, we have Professor Daniel Houser from George Mason. So, you'll have five minutes, and in the interest of staying on schedule, maybe we'll save questions for Pai-Ling until the end of the session, if we have time.

PRESENTATION DISCUSSANT -- PROFESSOR DANIEL HOUSER

~ Deleted From Transcript ~

DR. DURBIN: All right. So, next, we will have Luis Cabral telling us about the Dynamics of Seller Reputation: Evidence from eBay.

PRESENTATION: THE DYNAMICS OF SELLER REPUTATION: EVIDENCE FROM Ebay

BY PROFESSOR LUIS CABRAL

PROF. CABRAL: Okay, thank you very much. I'm going to be talking about some joint research I've been doing with Ali, who's sitting right here. So, if there's kind of difficult questions, I hope he'll be able to address those.

Now, Pat gave us a very good introduction to work on eBay and he characterized the various issues that we've been trying to look at. One of them is, what is the role of the feedback mechanism, what is the role of reputation on online auctions generally and in the particular care of eBay in particular? And that's what we tried to address in this research.

What I wanted to do first is, notwithstanding the very good introduction that we got from Pat, give you a little bit of what I think is the summary characterization of what the literature on addressing this particular question has been and how we have been trying to -- are trying to contribute to it.

I think that it's fair to say that when it comes to actual data -- so I'm not referring to experiments in here, but when you look at data, there are four different ways that you can get at it, and the four have been exemplified by a variety of papers that have been written on eBay and the online markets.

One is to look for natural experiments. The other one is to look for field experiments, and David and Pat and other people have been doing quite a lot of those. But then I think it's fair to say that at least in the particular case of studies on the value of reputation, most of the work that we've seen is on cross-section data, basically, you know, I just collect data on all the sales on a particular day or in a particular period of time for a particular type of object and then I try to do some sort of statistical regression where my dependent variable is the price that the seller was able to get or number of bids or whatever measure of success you want to consider. And then the explanatory variables are a variety of control variables, including some measure of reputation by the seller. So, that's sort of perhaps the most common way of assessing a valid reputation, and the answer is it's positive.

I mean, there are a variety of estimates that we all have done, but I guess the summary of it is that it's positive.

So, what is it that we have to add in here, you might ask, because we're not the first nor the second nor the third paper to look at this question. There's already a good dozen of them. And, in fact, most of them are very well summarized and surveyed in a paper by Ali and Pat in the Journal of Economic Literature.

What we're doing in here, which we think is, to a great extent, novel, is to look at the panel data of sellers as opposed to a cross-section. In other words, it's a data set that looks at the history of sellers, not just a series of sellers at a particular point in time. The advantage of doing that is that you can control for things that you might otherwise not be able to control.

That is, when you have a cross-section of a bunch of sellers at a given moment in time, even though it's true that what we see, sellers also see, there's still a variety of details about each particular seller that we may not be able to measure very correctly and, in fact, Pai-Ling was giving us a series of examples, you know, the way things are written, is it well-explained?

One of the most amusing sources of heterogeneity amongst sellers, perhaps not the most important, but in my opinion, one of the most amusing ones is, of course, whether sellers can spell in English or not.

(Laughter.)

PROF. CABRAL: There's a very interesting story in the New York Times recently about the value of our (inaudible). I mean, people will try to find items that are misspelled, like laptop with L-A-B, and they buy them because, of course, there are going to be very few bidders because very few people are going to be searching for those objects and then sell them at a more expensive price, spelling it correctly this time, of course.

(Laughter.)

PROF. CABRAL: That's the kind of thing that we, as researchers, could find out if we had the time to go and read through all of the descriptions in great detail. But if we don't have the time, you're probably going to miss that and there's probably going to be some error in the price of that laptop that was initially misspelled.

The advantage of having a panel data is that you're going to be able to correct for all of that, assuming -- that's the assumption that I would need in here -- is that if you don't know how to spell in English, you're going to misspell on a systematic basis every auction that you have. So, that's one advantage of the approach that we're taking.

The other one is that whereas most of the studies that we've seen on reputation on eBay focus on answering this first question, what is the impact of reputation on how much buyers are willing to pay, we believe that a related interesting question is, how does the system lead sellers to behave perhaps opportunistically or how do sellers react to the system itself? Do they make more effort? How do they change the amount of effort that they put through their life? Do they have any incentive to change their identity or to restart their reputation? I mean, there's a variety of, I believe, very interesting questions having to do with the feedback reputation mechanism that go beyond, do buyers react to it, which I believe is a very important one, but not the only one. So, those are the two ways in which I think we're trying to push things forward a little bit in here.

So, the summary of what we do is that we use a panel of eBay sellers using feedback histories as a proxy for their history. This is going to be where the rabbit is going to go into the hat, by the way, because, you know, no, we didn't wait for six years to collect our data. What we did was -- there's a certain amount of cheating in here, so I'm going to be very upfront about that -- is we're using the history of feedback, which is something that is accessible to us as a proxy for what the sellers have done in the past.

So, we're going to use feedback as a proxy for buyer satisfaction, but also the feedback rate as a proxy for sales rate and the absence of feedback as a proxy for exit. So, throughout my next 10 minutes, if I say that a seller exited, we define that as we've never seen that seller again trade on eBay, again meaning in the year-and-a-half since we started looking at this.

And the stylized facts that we find are that upon the first negative feedback, the sales growth rate drops dramatically for a seller, and more subsequent negative feedback arrives much faster than the first feedback -- negative feedback arrived. So, it seems like it's sort of like opening a gate for negative feedback when you get the first one.

We see that the exit rate -- again, exit defined in the way that I just mentioned earlier -- is declining with age and it's typically proceeded by a series of negative feedbacks. So, most negative feedback that a seller has are kind of concentrated towards the end of his or her lifetime, or at least his or her lifetime with that particular user name.

We also find interestingly that most sellers start off as buyers. So, if you look at the life cycle of a trader, he or she will start off most likely as a buyer and then become a seller. So, you can think of this as a form of building a reputation. Although, for my presentation today, I will not comment a lot on what is the theoretical interpretation of what the facts are, I will focus primarily on just describing facts, or what we think are facts.

So, I'm going to skip this part, in particular, because Chris, the discussant, said he would talk a little bit about theory. So, I think there's a natural complementarity because my presentation, talking about facts, and his presentation, talking about possible interpretations. So, we collected data in 2002, 2003. We looked at objects that we thought --

(End of Tape 1, Side A)

PROF. CABRAL: -- sufficiently homogeneous because we wanted to focus on the impact of reputation. Therefore, you want to abstract from variation in product quality. So, collectible coins, mint quality. A particular IBM notebook for which we knew there were relatively few variations in terms of the versions that would be a potentially problematic source of variation, you know, does it have more memory or less memory? This is a fairly uniform one. And the 1998 Holiday Teddy Beanie Baby from the Ty Company which, again, it's a very, very specific product.

We have about 130 unique sellers per product. Average feedback score is about 1,600. Average percentage of negatives -- this is true throughout eBay, if you're familiar with eBay -- is relatively small. So, a little less than 1 percent of negative feedback. And as you can see, we chose objects that had a certain variation in price, from the Beanie Baby, $11 on average, a relatively cheap item, to the ThinkPad, almost $600 of average value. So, that also allowed us to see whether there are big differences depending on what the value of the object is.

I think I'm going to skip summary statistics unless you want to go back to this. I gave you a summary of it. Again, I'm going to, very briefly, go through this slide. It has the distribution of feedback aggregate. This is something that you'll find, also, in a variety of other products on eBay that, as I said, the average number of negative feedback is relatively small. It's a very skewed distribution.

There's a good number of sellers that have a perfect record. In fact, quite -- I don't know how many percent they are, but it's a relatively significant number that have a perfect record.

So, let me tell you a little bit about stylized facts that we get out of this data. The first thing that we look at is, what is the impact of the first negative feedback that a seller ever gets? Since the number of feedbacks is so small, it's kind of a -- it's an interesting event study, as it were, what happens when your perfect record goes away? So, what we do is we look at the growth rate in the seller's sale growth rate from the week before the first negative to the week after the first negative, and then we would do the same thing for the second, third, fourth and fifth negatives.

And surprisingly, or perhaps not surprisingly, what we find is that around the first negative, the growth rate typically goes from a positive to a negative value and, therefore, there's a big difference, which is fairly significant. This is the standard error of the difference. And so, you know, for those of you not familiar with statistics, basically, what I'm trying to say is that this is a very precise estimate, a fairly precise estimate, and with a probability -- a P value of less than 1 percent.

Whereas for the second, third and fourth and fifth negative, which are not here on this table, we do not find, not even close, such an important effect on the sales growth rate. So, it's kind of the first interesting fact. The first negative seems to have a very big impact. Subsequent negatives, much less of an impact on the sales rate.

The second, interesting and somewhat puzzling, stylized fact is the frequency or arrival of that negative feedback. So, what we did in here, in this table, is we defined by T1 the number of transactions as we can measure them based on feedback, number of transactions before you finally get a negative. How long does it take for you to get the first negative? And then T2 is, how long does it take for you to get the second one after the first one? So, T2 is the difference between first and second negative.

Now, ET is not an extraterrestrial in here. ET is simply if you're to look at the lifetime of the seller and see how many transactions that seller had, how many negatives he had, were they equally distributed, what would be the frequency of negatives? So, ET is just the expected value of that frequency if they were equally distributed.

And here are the values. I mean, for all categories then, we can also separate that according to each of the four products that we consider. About 240 transactions until you get one negative, but it only takes 188 between the first and the second. And, in fact, if you are to look at the ET value, it's also low, in fact, even lower, 162. And this difference is relatively significant. So, the P value, if you will, the probability value for a hypothesis that T1 is greater than T2 is a little over 2 percent. So, it's relatively significant.

So, in words, what we see in here is that it takes a while before you get a first negative feedback. Once you get that first negative feedback, you're going to start getting negative feedback with a higher frequency.

The other thing that we find, it's not clear in here in this slide, is that this jump in probability then kind of flattens out. That frequency, more or less, flattens out after the second negative item. The third, fourth and fifth negatives seem to be arriving at a constant frequency. So, again, it's the first negative feedback that seems to have a big impact in here.

I don't have a lot of time, so I'm going to skip through these sort of more statistical uses of correcting for selection bias because there are potential problems in here with the sample that we're using. But that's in the paper that's been distributed. I'm also going to go over a variety of interpretations that there might be of why this feedback rate is going up so fast. We'll try a variety of explanations of whether it could be buyer behavior does it, and to cut a long story short, the answer is no. We are fairly convinced that this is not about buyer behavior. It's not about fear of retaliation. It's not about being afraid to be the first one to give a negative feedback. It's not about the first negative being a more negative-negative. We can strike the indices of nastiness of negatives and things like that, and that doesn't seem to be what's going on.

It's not about conformism and (inaudible), which is another possible interpretation. I think we've done a reasonable job at convincing ourselves that the fact that negatives are coming faster is not because buyers are changing their behavior and giving feedback. So, we believe it's something to do with the seller himself or herself.

I'm going to skip through this because I don't have a lot of time.

Let me talk to you a little bit about the other type of results that we have in here, which are the ones related to exit, again exit defined in the way that I mentioned earlier. A seller exited if that seller has not been observed trading on eBay again. That seller could very well be selling under a different name or what have you. But as far as we're concerned, that seller exited. In fact, as far as buyers are concerned, that seller has exited.

We see that a seller that has more positives is less likely to exit. That's what this regression says. And a seller -- yeah, and that's true for a variety of sub-samples that we considered. We also see that a seller who had more negatives, who had a worse record, was also more likely to exit opportunistically. I probably should say, what do we mean by opportunistic exit? And I should probably put opportunistic in quotation marks. An opportunistic exit is an exit that's preceded by a lot of negatives. So, if we observe that there's a seller with a lot of negatives in the last 25 transactions, we call that an opportunistic exiter. Quotation marks because it doesn't need necessarily to be opportunistic. It could simply be there are a variety of interpretations.

Be that as it may, consider that an opportunistic exit. Well, that fact that you, before that, had a bad reputation, that's a predictor of whether you're going to do that or not. So, if you have a record with a relatively high number of negative feedback, you are more likely to exit in this way, exit by accumulating a lot of negatives during the last few transactions that you had. So, another interesting fact.

And, finally -- and I think my time is running out -- we'll also see, as I mentioned, that many, many sellers have switched from being buyers to being sellers, and we then look at whether we can predict whether a seller does that or not, and it appears that sellers with a better record, in fact, are more likely to build a reputation in this way, which is an interesting thing -- point -- from the point of view of theory. (Inaudible) who wants -- who are the kind of sellers who would be willing to build their reputation in this way, if you believe that this is a form of building a reputation?

Finally, just to give you an idea of the order of magnitude of our facts, for example, the probability of exit in our sample is 18 percent. How important are positives and negatives in here? And one way of seeing that is to see -- and the -- by the way, the probability of opportunistic as it is 5 percent, according to our definition, which is a fairly high probability, I believe.

As you change the -- for example, the percentage of negatives from the 25th to the 50th percentile or from the 50th to the 75th, you know, here's how that probability changes, and that compares to an average of five. So, we're talking about big effects. I mean, these are not footnote kind of things, these are big effects.

So, there are a variety of theoretical interpretations for what's going on in here. This is not necessarily opportunistic exit. The frequency of negative feedbacks -- again, also, there's a variety of interpretations of why a seller might be increasing the frequency of negative feedback. We talked a little bit about that in the paper. We went from being very opportunistic to being very pessimistic to, I think, now being realists about how much we can say about this. I think there are interesting facts in here and I think this is kind of, at the very least, suggests that it may be interesting to look into it in greater detail, both from an empirical and a theoretical point of view. Thank you.

DR. DURBIN: To discuss the paper now, we'll have Chris Dellarocas of the University of Maryland.

PRESENTATION DISCUSSANT -- PROFESSOR CHRIS DELLAROCAS

PROF. DELLAROCAS: Thank you. Good morning, everybody. I've been familiar with this paper for quite a while and fascinated by the empirical findings. So, I hope that the authors will take well some of the comments that I'm going to make.

Let me first give you a summary of the empirical findings. So, actually, I will focus on number two and number three, which are the ones that Luis has also focused on. One of the main findings is that the first negative is actually creating a cascade of more negatives. So, negatives beget more negatives.

And another interesting finding is that just before exiting, sellers seem to accumulate a disproportionate number of negative comments. So, both of these findings are quite fascinating and I would like to focus a little bit more on what do these tell us, if anything.

Now, negatives seem to beget more negatives. It's very interesting. I mean, I was thinking about it for at least the last two, three years, which is how long this paper has been in gestation, and the main question here, in my mind, is, is this because sellers slack after they get a negative or because buyers stone them and just give them more negatives?

There are actually very good theoretical arguments for both explanations. I mean, the models for both moral hazard and adverse selection can explain seller slacking after receipt of the negative, and then we're going to have models of conformism that can give some explanation as to why buyers seem to be stoning.

One aspect of eBay that is not discussed in the paper and I would like to add is that eBay's feedback mechanism is very vague. EBay doesn't give you clear guidelines as to what exactly you're supposed to rate and under which circumstances you're supposed to give a negative. So, it's to be expected that some sellers have uncertainty and they might be looking at what other people have done before them for guidance. So, you can actually form a very good theory as to why buyers might stone sellers after the first negative.

Now, Luis claimed in his paper with Ali that they believe it's slacking and not the stoning, and I'm not entirely convinced of most of their arguments. To their credit, they put together a lot of arguments, but most of them are informal.

Now, there is a very recent paper of Kopker, Lee and Reznik that is looking at exactly the same thing, and they did something which I think is actually quite clever. They compared the incidence of negatives following the first negative for auctions that have completed after the receipt of the negative against auctions that have completed before the negative was received, but received feedback afterwards.

Based on this analysis, in that paper, the authors were able to find some strong evidence of -- actually, I would say stronger evidence of stoning rather than slacking. So, their tentative conclusion, again, is what seems to be happening is not so much that sellers change their behaviors, but buyers are more likely to stone them after the first negative.

Again, I think that this is still open. I mean, no paper has actually given the final word on this and it's very interesting to see what more we can do to figure out what's really happening in this case. And, of course, there are many interpretations other than those two, like there might be a grace period, like, for example, buyers can be very lenient until you get your first negative and then they might shift to another mode of behavior, which is -- and that might be the normal mode of behavior after that.

Then negative feedback and exit, again, the causality is not clear. Is it that sellers decide to exit and then they milk their reputation by cheating the buyers or is it that sellers are unlikely -- they get a stream of bad feedback and then they get discouraged or they decide that they're better off to disappear and reappear with a new identity, and that's why you see the effect. I mean, the empirical results will be identical in both cases.

Now, in terms of the interpretation, actually, let me very quickly say, the empirical assumption that underlies this study is that the rate of feedback submission is constant. So, they use feedback as a proxy for sales.

Now, we have a number of empirical studies, including one of mine and my results show that this is not the case, that the rate of feedback depends on a lot of factors, such as the buyer and the seller's reputation score, and also that paper, as I mentioned by Kopker, et al., seems to find that after the first negative, the rate of both positive and negative feedbacks decline. So, after you get a negative, people are reluctant to give you both a positive feedback and perhaps this might bias the empirical rate of incidence of a negative feedback.

Now, in terms of the theoretical models, again, the authors, to their credit, they don't make an upfront assumption that this is the model and not data supported, but they actually go through a number of models and they discuss to what extent the data supports the models. What's interesting is that under certain assumption, the empirical results that they observe can be explained by either a pure moral hazard or pure adverse selection or a combination of moral hazard and adverse selection. So, it is not crisp what exactly is going on here.

One thing that struck me, however, as I was reading the paper is that all of the models -- in all of the models, the behavior of the seller deteriorates in a way after the negative comment. Somehow, to me, this doesn't seem right. Intuitively, there's something that makes me very uncomfortable with this. I mean, what about learning? I mean, one would assume that the negatives are also a learning experience. In our life, if we do something wrong, I mean, we learn and we become better, and we don't seem to see this in this data.

So, I didn't feel uncomfortable with the fact that there's no learning on eBay and sellers don't improve. So, what I'm beginning to suspect is that what's really going on is that there's a lot of different segments within eBay that are reacting differently to feedback. So, for example, there could be -- let's say, the results of this that the office observed can be due to the fact that, well, buyers stone, okay, so, you know, after your first negative, they are more eager to give you another negative, and when it comes to sellers, well, half of them slack and the other half improve. So, seller behavior cancels itself out, and what we see is primarily due to the buyer behavior and not the seller behavior.

And, of course, each of you can give a number of different interpretations to what's really going on, which leads me to my conclusion. I think this paper is actually very thought-provoking, not only in terms of the results themselves, but in terms of what it might mean for methodology and further research. We have a paper where the findings are fascinating, but where the authors, despite doing a very, very thorough job of considering a number of different interpretations, they're not able to give a crisp conclusion as to what's really going on.

And what -- I mean, one hypothesis that I don't think they considered that might be plausible is that, well, all of these phenomena are actually taking place. We have a huge market. I mean, eBay has a very, very large scale market, and there might be -- and, also, eBay's feedback mechanism, as I repeated before, doesn't give you clear guidelines. I mean, it doesn't really -- it's a very vague mechanism. It doesn't give you clear guidelines on how to interpret the information and how to react to it. So, it might be that different segments are actually reacting to the information in different ways.

All of this phenomena are happening at the same time which, of course, is a very interesting challenge for researchers. Because -- which means that, you know, we have all this data on the Internet and, yeah, we can get it and, yeah, we can get a lot of interesting population level correlations, but then if you want to mark those population level correlations to individual level inferences, it's very difficult. Either one assumes a theory and then you say, gee, my data has -- fits the theory. But then, again, there are 10 more theories that fit the same data.

I mean, so, I really think that the authors did a very, very honest and good job of not really sticking with one theory but considering several of them. But then, if there are many theories that are consistent, what happens? So, I think this is a very interesting question that I would like to finish with. I don't really know what the answer is. One answer would be to experiment with mixture models, very difficult to estimate, but, you know, perhaps that's something that we should be taking more seriously and, again, another answer is -- I would like to resonate with the previous discussion -- to supplement the data we collect from the Internet with data we collect from different methodologies, such as surveys or controlled experiments.

So, I really like this paper because it -- not only did it give us some window into what's happening on eBay, but it really makes us think about all of this very interesting methodological issues.

DR. DURBIN: All right, thank you very much. So, again, in the interest of time, I think we'll push on to the next speaker. So, next, we'll have Ali Hortacsu from the University of Chicago tell us about the geography of trade on eBay.

PRESENTATION: ON THE GEOGRAPHY OF TRADE ON Ebay

BY PROFESSOR ALI HORTACSU

PROF. HORTACSU: Thanks a lot, Eric, and I, first of all, thank Chris for his extremely valuable and insightful comments. We're going to talk -- we should talk afterwards.

So, this is a very new paper, particularly compared to the other papers that, as Chris said, have been around for around two years, maybe more than two years. It's a very recent paper, so please, if you have any comments, any reactions, if you think this is completely wrong, you know, please come tell me because I cannot explain some of these findings myself, and this is joint work with Jason Douglas and Asis Martinez-Jerez.

I guess the questions that moderated this exercise came from a book titled Death of Distance by Cairncross. This book describes sort of what will happen with the invention of the Internet if you lower search costs and buyers are able to find sellers everywhere. The first paper I know that actually tries to look at this, is by Austan Goolsbee, who looks at people's purchase behavior online. He finds that people who live in high sales tax states are more likely to purchase online to avoid the sales tax. That's his paper in 1999 in QJE.

And then there is another paper that finds some correlations as to cross country trade volumes and the number of Internet hosts in this country. There’s no causal interpretation of this finding, but it's an interesting correlation.

Also, there's a paper by Avi Goldfarb, which is probably the closest to what we're going to do. It looks at the clicking patterns of a sample of web users and whether people look at other countries’ websites. Basically, the conclusion is that on the Internet, local content matters. If I'm a Korean speaker, I'm going to look at Korean newspapers and maybe not others, or non-Koreans will not look at Korean newspaper's findings.

So, what are we going to do? Well, I should also say a bit about some of the literature in the trade literature. Researchers in international trade find a huge amount is explained by distance and the size of the trading economies. The big stylized point from the international trade literature is that if two countries are far apart, the amount of trade between them is less. The empirical finding is that something called the gravity equation holds.

So, what is the gravity equation? If you remember high school physics, think about the gravitational force of attraction between two bodies. The gravitational force of attraction is the product of the masses of those two things divided by the distance between them squared. Take the log of this and you're going to have basically the log amount of gravitation attraction is the log of the mass, either together and some distance squared.

In the trade content, gravitational force is the amount of trade between two countries and the masses are the sizes of the countries. And, again, you have the distance being a factor. There have been hundreds of papers that find that this equation fits trade patterns very well, and that this distance matters quite a bit. Among distance crossings, it's not just the physical distance, but also crossing a border that enters as a distance coefficient.

So, what is this exercise? This exercise is to run the gravity equation to see if the gravity equation holds on eBay. And if it does, what are the patterns it reveals?

I should not have to tell you about why eBay matters. EBay matters a lot. In a recent finding, apparently 30 percent of households surveyed in 2004 had bid on online auctions. Stock market participation rate is maybe 50 to 60 percent. This is as big a household phenomenon as it gets.

So, why eBay? Why look at eBay? Different people have different priors as to whether the gravity equation should hold and whether distance should matter on eBay. My prior coming into this research was, why should it matter? Why? First of all, there are no search costs. So, one big factor that might prohibit me from buying from somebody far away is that I am not aware of this person's existence.

Another thing is that for many items on eBay, shipping costs are uniform in the sense that I just put it in a package prior to mailing the envelope and sending it off. There are the postal zones, but for a lot of the items, the postal zone rates do not make that big of a difference. At least, the postal rate zones are much larger than, say, the sort of zone for the distances considered here.

Another thing to consider is uniform market mechanism. You don't have different institutions, it's just one institution that governs everything.

This was my prior. Other people's priors were, well, how did eBay start this? EBay started as a substitute for classified ads to the newspapers or the market for used goods, and these markets were very local on character. For instance, you would look at your local newspaper if you were just moving into the town or trying to buy something. This might suggest a much more local feature of the market.

So, what did we do? Here, we wanted to get some breadth, so we sampled the main categories -- the 30 main categories on eBay. I think, actually, it's 29 because we don't have data on eBay Motors, which would be interesting to look at.

So, we got a random sampling of listings in each category. Conveniently, eBay shows the buyer’s location, if the buyer enters the information. Now, correct me if I'm wrong, I believe eBay actually provides seller's information based on registration information also. But our experience was that this information was not always available in the sample that we looked at, which was only 2004, and sometimes very vague. So, there was a judgment call as to where the seller was located.

But, actually, the biggest caveat with our data set is that we had to find where the buyer was located. We could only see the buyer's location if the buyer also made a sale. If anybody is willing to help out with buyer locations, this would really add to the study. One thing in particular that we see is that the buyers whose location information we could find by looking at their sales are much more experienced than the buyers for whom we couldn't find location information.

I won't spend much time on the model. The model is a simple one with different types of sellers and different types of buyers. Here, type means location.

Basically, all of the buyers are heterogenous, except that type and location are going to come through some error term, and so, people have some random distribution evaluation for an item. Aside from that random distribution evaluation, it's the distance that matters as captured here by the interaction between the buyer and seller location.

The main assumption is that the auction mechanism is efficient, which remains to be proven. However, I suspect this is probably the case with so many auctions going on together. And what does efficiency mean? It means that it will award each good to a buyer of the highest willingness to pay. So, if this is true, then controlling for the distance, the person with the highest valuation for the item should win the auction. The probability that a buyer located in Location B wins the object of a person located in Location S should be given by this expression, which will look very familiar to those who have dealt with larger demand systems.

This was all an exercise to try to motivate a demand regression. I'm going to put the log number of transactions between buyers in Location B and sellers in Location S on the distance between the two buyers, between the locations, controlling for things like market size. Well, market size indicates number of sellers or number of buyers located in the market. In most of our analysis, we're going to put in buyer and seller location fixed effects. So, it is going to take out a lot of the heterogeneity.

Again, this is just Newton's law of gravitation equation where you're looking at the log amount of trade, controlling for market size on both buyer and seller markets, and putting in distance as an explanatory variable. You can think of it as what explains the volume of trade controlling for the size of the markets, and how does this depend on distance?

This table might be a little bit small, but this is just to show how our first set of results compared to what other people have found in trade studies. So, let's focus on the distance coefficient here. In two studies Wolf and Hummels, these last two columns, were based on actual trade data from consumer survey. They found a coefficient of minus one on eBay. Well, to verify my prior, the distance effects seem much smaller, about a tenth the size of that coefficient, although something interesting happened when we first ran these regressions.

In both Wolf and Hummels, they found a big coefficient on an indicator for the transaction happening within the same state. So, this means controlling for distance, there's more likely to be transactions within that same state as opposed to another place that's equally far away. And our coefficient size is similar to what Wolf and Hummels finds, even though the distance coefficient is much smaller.

So, then we wanted to explore this a bit further. We said, well, what is it going on in the same state? Is it sales taxes, is it some state laws? But we had the good fortune of having a formation more detailed than the state of the seller. So, I will show that.

Next, what we had was data on the city locations of the sellers and buyers in certain cases. This reduces the amount of the data we have, but when you do the same regression using the city information, the picture emerges as the following. This horizontal axis is the log of the distance in kilometers. So, this is basically the same city. This is people located 25 kilometers, people located 50 kilometers, 100 kilometers, 250, up to 4,000 kilometers. Basically what this says is there seems to be a big bias towards transactions within the same city or within driving distance. But beyond driving distance, distance doesn't matter at all.

eBay seems to be aware of this in the sense they have a search interface that allows you to look at items within your zip code. They have the local search option, and I've used this actually a couple of times. If you look through the chat boards, people were very peeved when eBay took this under maintenance. But when eBay put improved features on the local search option, they appeared to be very happy.

But what we did next is to try to understand why there's a same city bias. So, we estimated this same city bias coefficient across item categories.

PROF. HORTACSU: To explain this bias, if you were a sports fan, you would buy stuff from your own city.

So, local preferences matter. Let me show you some graphs when we take out those first two categories from this table. This the first graph. This is where we'll plot that same city bias against average object value in category. I could fit a straight line through these, and it seems like the more valuable the item is, the bigger the same city bias is. You're more likely to buy it from somebody within the same city if the item is more valuable. Jewelry and watches don't seem like particularly hard items to ship if you're not going to insure it, but business and industrial objects might be heavier.

Here, we looked at the seller's reputation profiles in each category. We sampled some sellers from each category, and we looked at percent of negatives in seller's profile. It seems like consumer electronics is the big outlier here with the seller having above 1 percent negative profile. Other sellers seem to be less negative-prone. Still, here, there seems to be a correlation between the badness of the seller, or the number of complaints against the seller, and this bias towards buying from someone from the same city.

Here's the regression that we ran. We sampled items from these categories, and we assessed the weight of these objects to see if weight is determining. Because if I'm going to buy a car, I'm more likely to buy it from a person in the same city because it's heavy.

It seems like the percent negatives in the seller's record and the value of the object seem to come in much more significantly than the weight of the object. That may explain quite a bit of variation here.

So, I want to leave with a bit of a plea for some answers, because the story that comes to mind after this is some issue of trust. You want to actually pick up the item or you want to be able to return it to somebody. I would have expected this local pick-up to be much more important for things like couches, but this didn't seem to be the case.

Another alternative explanation that I've been thinking about, which might be of interest to the people from eBay, is whether this is a manifestation of shill bidding. Shill bidding is basically the seller buying his item himself or trying to inflate the price by bidding on the item himself. But then, why would the seller, who's sophisticated enough to do this strategy, declare his location truthfully? Buyers are aware of this possibly.

I should also point out that our data set shows the buyer's location only if the buyer made a sale. So, I cannot think that sellers would use shilling identities, actually, because you want to concentrate your sales on items for which you have the most feedback on. But I think this is an interesting hypothesis to study, because if this indicates shill bidding, it might be a very strong indicator of it.

So, I just want to conclude, why isn't business dead? It seems to be dead beyond the city limits, but it's not dead within the city limits because, yes, there are local preferences. Mainly, there's a trust issue. I trust people who I can pick up stuff from more than other people. Some people might say it's an immediacy bias where you have to wait for something to be shipped rather than driving. I would have expected somewhat more of a linear response to that. The other question is whether this is a shill bidding manifestation or not. So, I should end with that.

DR. DURBIN: Thanks. All right, so, for discussion, we'll have Robert McMillan of the FTC.

PRESENTATION DISCUSSANT: DR. ROBERT McMILLAN

DR. McMILLAN: Thanks.

(End of Tape 1, Side B)

DR. McMILLAN: -- and I thought that it was very fascinating, the extent to which most of the effects occur within the driving distance. There were a couple of things that I thought would be more interesting to look into, going forward. The first is to delve deeper into looking at this substantial increase in the willingness to pay within the driving distance. You sort of looked at this a little bit at the category level using the data that you already have trying to sort out trust versus immediacy versus the correlated preferences of the buyer and the seller – the reason that the ticket prices tend to have a higher willingness to pay within driving distance.

Also, there may be shipping costs. If you're trying to buy a computer monitor, you may have a much higher willingness to pay if you're located in the same city – you're willing to drive across town and your shipping costs are pretty much zero – whereas if you wanted to actually ship this item it would be expensize – it's a bulky, heavy item. So, your willingness to pay would be higher.

Going forward, I would urge you to look more deeply at those things. I know a lot of the tables you had were at the category level, and looking at seller ratings at the category level is a really rough measure, but I think you're off to a strong start.

The second thing that fascinated me was the increased willingness to pay for tickets. I can think of two simple explanations for this. The first is that, to a certain extent, tickets are "experience" goods. So the relevant shipping cost in this case is not the cost of shipping the ticket – you just put it in an envelope and send it through the mail – it's the cost of shipping the person from wherever they live to the event site. And that's a very bulky, typically very heavy object.

(Laughter.)

DR. McMILLAN: So, that's one part of it. The other part is the correlation of preferences. I think correlated preferences could be really interesting because a lot of what's neat about geography is looking at boundaries, and I think you could use tickets to sort of sketch out relevant boundaries around cities with an increased willingness to pay for tickets. So, using the correlation of preferences for, sporting teams, you could, in theory, sketch out a boundary of where the fans are for that particular team.

I have a few more specific comments that I can give you offline because it's a little hard to get into the specifics here discussing the regressions. But, anyway, thanks for a really interesting paper.

DR. DURBIN: All right. Well, I think we've just about exhausted our time for this session, so why don't we go straight on to the coffee break here. We'll have 10 minutes and start back in a couple minutes after 11:00. So, if you have questions for the speakers, you can find them drinking coffee.

(Whereupon, a brief recess was taken.)

DR. DURBIN: All right, if people can take their seats so we'll get started again here.

(Whereupon, there was a brief pause in the proceedings.)

PANEL: FRAUD, INFORMATION AND REPUTATION

DR. DURBIN: All right, so for the next session, we have a panel which will consist of a whirlwind set of presentations. We'll have 10 minutes for each speaker and then, hopefully, a little bit of time for questions at the end. So, without further adieu, to get started, Ginger Jin is going to talk to us, I think, about baseball cards on eBay.

PANEL PRESENTATION BY PROFESSOR GINGER JIN

PROF. JIN: Thank you so much for including me in this very informative conference. I want to make some comments on price, quality and reputation relationships on eBay. I would first go through what I have learned from my own research. Then I want to talk about how my research links with the studies that have been shown earlier today, what I have learned from those studies, what are the remaining questions, and what I think is the policy relevancy in this area.

Pat has made the point that eBay is a special auction market with a very large scale. I see it more similar with the traditional markets where people could have different information on different sides. Obviously, quality is not perfectly observable on eBay. We don't know exactly whether the seller would deliver and we don't know whether the delivered quality will be as good as promised.

There are a lot of theoretical literatures on the markets where quality is not observable. Most theories predict that, if we don't have an explicit indicator of quality, we may see a monotone relationship between advertising or reputation with quality and price. The logic is based on repeated sales where the buyers are able to tell the quality after purchase and they're able to track the seller's identity and talk to each other. In this framework, sellers of high quality goods get a lot of repeated sales, obtain good reputation and therefore are able to afford advertising. This supports the monotone relationship between price, quality, advertising and reputation.

The goal of my own study is to check this relationship in eBay. For every assumption in the reputation theory, we see some counterpart in eBay. But if you're familiar with eBay, you may realize that the eBay mechanism is not perfectly aligned with the theory. So the empirical question is, given the discrepancy between the reality and the theoretical assumptions, do we still observe this common classical insights applicable to eBay? Or where is the difference?

My own study has a very simple research design. We first watched the market for seven months, which tells us the probability of completing an auction, the final price, what a seller has claimed about quality, and what was the seller or buyer reputation in terms of eBay ratings. We looked at baseball card, like the one I am showing here. A real baseball card is a little bit bigger than a regular name card and it often depicts famous players from the baseball Major League. A baseball card can be very expensive if it's for a very famous player, if it's a rookie card, and if it is kept in good conditions. The one I'm showing here could go as high as $1,400 if it's in very good condition. In fairly good condition, it could still be as expensive as $100. So, we're not talking about pennies, we do talk about real money here.

One thing missing here is the information that sellers know but we consumers probably don't know, that is the true quality of the card. To obtain information on card quality, we conducted a field experiment for ungraded cards. The cards appeared as ungraded, so we didn't know the exact quality. We purchased the cards systematically so that one group has very high seller claims of quality and the other group has moderate seller claims or no seller claims. Then we sent them to professional grading, which told us the true quality. Now we have every information we have through the market watch plus the true quality. We're able to see exactly how much is the information missing and whether buyers are able to tell this information from what they have observed on eBay.

I'll go through the facts very quickly. If we just look at the market, we do see higher price for graded than for ungraded cards. This is pretty familiar with economists, only the good ones are going to be graded given there's a cost for grading. For ungraded cards, we see a significant price premium, about 20 to 50 percent more, on those who claim high quality. Note that the claim is immediately incredible. A claim of cards in gem or mint condition, if it's true, means that the card could be sold for as high as $1,400. But here the seller is willing to sell it by $100. Buyers certainly recognize that the quality is over-claimed, but the puzzle is, they still pay some significant amount for that claim.

Another fact is that seller reputation is positively related to the probability of completion, but conditioned on completed auctions, it has no impact on final price. Based on these facts, it seems roughly consistent with the theory -- we see the price going up, we see the advertising in terms of claims going up, we see the reputation going up, and then we may infer those transactions mean good quality.

But some puzzles cannot be explained by a simple market watch. For example, reputable sellers are less willing to make claims. Buyers with more ratings seem not willing to bid on graded cards, especially on those with high claims. These two facts are inconsistent with the classical theory, but we can explain them by our experiment.

The experiment basically shows two facts. The first is, those who made high claims are more likely to be fraudulent. By fraudulent, we mean the sellers don't deliver the card or they deliver a fake card. Conditioned on those who do deliver an authentic card, we see no better quality after delivery as compared to those who made moderate claims. This fact means that there's a very good indicator of who is a fraudulent seller. We can just look at who makes high claims. If people know the true relationship, they should bid lower on high-claimed cards.

Another fact is that reputable sellers are less likely to commit fraud. So they do send us something authentic. Conditioned on that they send us something, they don't send better quality. The quality seems unrelated to seller reputation. This means that, in some range -- I'm not saying in the whole market, in some range, there is a reverse relationship between price and quality, and in some range that we don't see a monotone relationship between reputation and quality.

What do we learn from this? Some facts are obviously inconsistent with the theory. The first thing I would like to stress is that we should be careful about market watch data, what they infer, and how they apply to the theory. From the surface, it seems like everything is perfectly aligned and the theory is right and somehow eBay is able to figure everything out. But the truth is quite different.

Another lesson is that, at least, some inexperienced buyers are misinformed. As Pat and Pai-Ling have mentioned, one explanation is that those who make high claims probably want to generate a huge information dispersion and, therefore, to generate a higher price according to the auction theory. But there's still a remaining question of why those claims are able to generate some buyer misconceptions. Why are some buyers willing to overestimate the value of the auctioned item? There must be a mechanism, either those buyers are naive enough to believe the incredible claims, or they're not familiar with the true facts in this industry, or something else.

As you know, the price-quality inconsistency opens doors for fraudulent sellers. Some people naively believe something, sellers are going to take advantage of that. This is very relevant to Luis and Ali, the seller dynamics. Actually, many facts we observe in our data are consistent with their findings. For example, sellers are more likely to start as a buyer and if they have negatives, they're going to exit the market.

One disturbing thing we find is that seller dynamics could be manipulated to facilitate fraud. We have seen people who defrauded us, set up their reputations through buying, struck a series of frauds in a very short window, and then, after they got a series of negatives, just abandoned the ID. In this way, one can conduct frauds very easily and does not get caught.

One remaining question is buyer dynamics. In our data, we see some evidence that buyers are learning over time. They get burned once, then they avoid this very troublesome market, or go to a safer place such as the graded market. What I really want to know is how fast buyers learn about this. Do they totally exit the market after learning or do they move to a safer market? How do they recognize the signals of safer markets across different categories in eBay? Answers to these questions are very likely a limiting force on fraudulent sellers. If the buyers could learn very quickly, then sellers would have a very short window to be able to strike frauds. If the buyers learn very slowly, then we could have a large room for fraudulent sellers to exist and probably prosper there.

This also raises a policy question. To what extent should we educate consumers? Shall we tell them the real risk of fraud on eBay? What signals indicate fraudulent behaviors on eBay? Who should take the responsibility of educating consumers? Should FTC do this or should eBay do this? Or should we just rely on sellers to give buyers whatever information they're willing to give on quality? These questions are very important to combat frauds, and, more generally, to speed up the solution of adverse selection problems in this market.

Another lesson I learned is that reputation does provide us some information. Many empirical studies have documented this, but I would say the evidence is still limited – it is not as much as the theory predicts. Here, we see the reputation signals delivery reliability, but they don't signal the true quality of the good.

This could be attributed to several settings in eBay. For example, eBay gives the same ratings no matter whether the transacted good is worth $1 or $10,000. The rating only shows the information on delivery but does not shows the value of the goods delivered or the quality of the delivered goods. We casually see some comments on the quality of delivered goods, but that could be subject to people's idiosyncratic evaluation, which is hard for newcomers to comment on. We also see sellers try to combat with some negative comments of their quality, by saying that the buyer was over-optimistic about what he has promised. These comments are difficult for outsiders to make a judgment of which side is really right.

We have also seen anonymous ID switch. In some examples, we see that people abandon old and bad ids. They don't care about reputation in these ids because buyers won't be able to track the same person across different IDs. This raises a second policy issue. When I say policy issues, I don't always mean public policy, this could be eBay policy or it could be other private sector policy. No matter who implements the policy, the question is, how should we improve the reputation system so that it becomes closer to what a theory would predict and gives better information to buyers?

When we have a better reputation system, we're going to increase the cost of fraudulent behaviors on eBay. For example, is it possible to track different IDs for the same seller? In a tracking like credit report, an individual goes with a true identity and that true identity will follow her regardless of how she changes her name. In a world that history follows, one should care about the history. This would strengthen the reputation system, and therefore reduce the number of frauds on eBay.

That's all I have. Thank you so much.

DR. DURBIN: Thank you. So, next I think we'll move to Keith Anderson of the FTC who will be telling us about --

MR. ANDERSON: How did I wind up next?

DR. DURBIN: You wound up next.

MR. ANDERSON: I'm supposed to be last.

DR. DURBIN: Well --

UNIDENTIFIED MALE: You got promoted.

PANEL PRESENTATION BY KEITH ANDERSON

MR. ANDERSON: It's a pleasure to be here today, but it also seems somewhat strange because I don't know anything about the subject of this conference. I've not done work on Internet auctions. I'm not even one of the 30 percent of people who have bought something in an auction. So, what I thought I would do is talk about something I know at least a little bit about: the database of consumer complaints that the FTC has and see what that can tell us about Internet auctions.

Before going forward, I guess I should issue the standard disclaimer. What I'm going to say represent my views and my views only. They don't belong to my management, they don't belong to the Commission or any Commissioner.

Oh, and before I go on, I'll make one comment on Ginger's paper. One of her findings was, if people make incredible claims, they often don't deliver. This seems very much like a message that the FTC tries to get out every day: If it's too good to be true, it probably isn't.

So, I'm going to talk about what we call our Consumer Sentinel database of consumer complaints. This is a database that collects complaints on consumer fraud and identity theft. It's maintained by the FTC, but it contains both complaints received directly by the FTC and complaints that are referred to us by over 150 other organizations who share their complaints with us. Sentinel was started in about 1997 and now contains about two million complaints.

As shown on this first slide, more than 635,000 complaints were received in calendar year 2004. Over 245,000 of these involved identity theft and almost 390,000 involved consumer fraud.

Looking now more directly at Internet auction fraud, Table 1 shows that Internet auctions were the subject of almost 100,000 of the 390,000 fraud complaints received in 2004. Internet auctions accounted, therefore, for almost a quarter of the fraud complaints in 2004, and the same was true in 2003. Indeed, Internet auction complaints have accounted for the largest share of fraud complaints received each year since 1999. The number of Internet auction complaints increased almost 20 percent between 2003 and 2004, and almost doubled between 2002 and 2004.

If you look at the figures for the first half to this year, it looks like things might be down a little bit, but that's just because it's the first half as opposed to the second half. If I actually compared the first half of this year to the first half of the year before, the number of complaints is up something like 20 percent, and as a percent of all complaints, it's just about on line with last year.

A word of caution here in terms of looking at the intertemporal comparisons. It's always a little bit risky in looking at our Consumer Sentinel data to look at the data for successive years and say, oh, problems are going up, problems are getting bigger, we're getting more and more complaints all the time. One of the things that the people in the Bureau of Consumer Protection who run this program do is to constantly try to get more agencies to report their data to us, to give us their data. So, some of this growth can be we're just getting a better picture of what's going on, not that real changes are going on.

So, with this caveat in mind, what I'd like to do is take a quick look at a couple of aspects of this data. First, I would like to look at some of the characteristics of these complaints that we're getting. First of all, Table 2 summarizes the types of problems being described in these internet auction complaints.

I should make at least one methodological point here. For anybody who is familiar with the sort of annual reports that the Commission puts out on its Consumer Sentinel database, the methodology here is somewhat different from what is used in generating those reports. When people file complaints, either with us or with some of our partners, in general there are a number of specific pieces of information that you're encouraged to provide. So, there are specific fields where you can provide this information. And then there is a field where you can also provide a textual description of the problem.

When the standard reports are done, they're all based on the information in the specific fields. On the other hand, what I've done is to look at the textual fields and try and describe what appears there. By reading the text fields, you can sometimes get more information. For example, people sometimes won't fill in the amount paid, but they'll tell you in the text that they sent a check for $200. So, the two methodologies are not directly comparable.

The second thing I guess I should note for clarity is that when I've got percentage figures here, it's the percent of the people for whom I could get the particular characteristic. So, I've not included the not reported in the base there.

Finally, I should also note that the figures I am presenting are based on a fairly small sample. It's supposed to be 125, but I made a minor mistake, so it turned out to be 129. And it's a random sample -- a reasonably random sample – of complaints received in calendar year 2004.

With these caveats, I note that almost three-quarters of these complaints deal with situations in which the person complaining says that they won the auction and paid for the goods, but never received them. The question of quality that Ginger was trying to address comes up in about 15 percent of cases. Other problems account for the other 10 percent of complaints.

In terms of the amount of money involved, the amount ranges from zero to $16,000. There was one case in which somebody filed a complaint even though they hadn't actually fallen victim to the fraud. However, they were trying to pin down some detail with the guy and he kept giving him run-arounds. So they felt he was probably lying to them. There was no money involved there. At the other end of the spectrum, there was one complaint that involved $16,000. I can't remember what the product was. The median amount involved in these complaints was just under $300.

Another thing that seemed to me might be interesting to look at was the payment mechanism involved in the transactions. Were we talking about payments being made in ways where the buyer had some recourse if they didn't get the goods or if there were problems, or were they paying with cash or other payment mechanisms that provided little recourse? It turns out that Paypal was the most frequently cited, being the mechanism in something just over 40 percent of the complaints. However, payment mechanisms that are much more cash-like, things like wiring money, money orders, checks or cash were cited in just over 50 percent of the complaints.

I should be a little cautious here. I don't have data on what percentage of all Internet auction transactions involved Paypal or these other mechanisms. These are the percentages of the complaints that involved these transactions. Therefore, you can't look at this and say, oh, well, Paypal is the most risky payment mechanism out there. It may be that Paypal is used in, say, 80 percent of transactions on eBay or in Internet auctions, and that would suggest that it's only half as risky as these other mechanisms. So, you've got to be careful about how you read this.

Another question that I was kind of interested in involved the amount of time that elapsed between the transaction and when the complaint was filed. After anecdotally looking at a few of the complaints, it seemed like the time that was elapsing between the transaction date and when the complaint was being filed was awfully short. And I was worried that people may have been complaining before there was really enough time to find out whether they were really going to get the goods. So, Table 5 looks at that question. As you see there, 10 out of 83 complaints – or about 12 percent – indicated that less than 10 days had elapsed. I leave it to others to decide whether that tells us anything or not.

Let me now turn real quickly to a second issue that I think is worth addressing. If we go back to Table 1, which I'm not going to do, because I'm not that fast with this equipment, you'd see again that 25 percent of the complaints about fraud that we get involve internet auctions. It's the biggest topic year in and year out.

So, it seems fair to ask whether this really says that internet auctions are the biggest problem out there. Are Internet auctions really responsible for 25 percent of fraud?

Tables 6 and 7 present one piece of evidence that is consistent with the notion that all internet problems, not just internet auctions, may be over-represented among the complaints we receive. One of the questions asked of consumers who file fraud complaints is how did the company that is the subject of your complaint first contact you? If you look at Table 6, and particularly the first two rows of that table, you'll see that for each year, 2002, 2003, and 2004, more than 50 percent of those who answered this question indicated that the initial contact was on-line, either e-mail or a website.

Compare this with the findings of the survey that we did a couple years ago, in which we asked a random sample of consumers about their experiences with consumer fraud. This survey was actually conducted in 2003; we're doing it again in about a month. As part of the survey, those who had experienced consumer fraud were asked how they first learned about the product or service involved in the fraud. As shown in Table 7, only 14 percent of the victims who answered this question indicated that the initial contact was through internet or e-mail.

Differences in the wording of the questions, etc., may be responsible for some of this difference. However, at least as a first cut, it certainly seems to suggest that complaints are more likely to be generated in on-line transactions than in fraudulent transactions in general.

Looking quickly just at internet auction problems, rather than all the transactions, it occurs to me that one reason we might expect complaints about auctions to be over-represented is that the companies involved in the internet auction process – eBay and BidPay are the ones that I'm specifically aware of though there may well be others – make it particularly easy for victims to file complaints that find their way to the Consumer Sentinel database. These firms provide direct links on their websites to on-line complaint sites maintained by the Internet Crimes Complaint Center, also known as IC3, which is an on-line complaint center maintained by the FBI, or to websites maintained by the FTC.

I would note here that more than 90 percent of the internet auction complaints received in Consumer Sentinel during 2004 dealing with internet auctions were filed either with IC3 or on-line with us. In contrast, the percentage of complaints coming from these two on-line sources is much lower when we look at some of the other frauds that the FTC worries about. With other frauds, the percentage coming from IC3 or our online complaint forms can run in the 10 to 15 percent range. Again, this doesn't prove anything is going on here. However, the results are certainly consistent with the fact that we've made it real easy for people to complain about Internet auction problems and they are doing so.

Let me be clear here, and then I'll stop. I'm not suggesting that eBay ought to make it harder for people to complain. Making it easy to complain and therefore having a larger sample of the universe of problems gives us a better picture for what's going on. When our lawyers want to take law enforcement actions or when I want to do the kind of analysis that I've done here, I get a better picture.

But if it's easier and less costly to complain about internet auctions than it is to complain about some of the other frauds that the Commission is concerned about, you've got to be real careful in interpreting the relative number of complaints for the two groups.

So, I think I'll stop there.

DR. DURBIN: All right, so next we'll have Debra Matties.

PANEL PRESENTATION BY DEBORAH MATTIES

MS. MATTIES: Hi, My name is Debbie Matties, and I'm an attorney in the Bureau of Consumer Protection at the Federal Trade Commission. I suspect there are some economists here from the Bureau of Economics, so much of the subject matter I'm going to cover in my slide presentation at the beginning is probably known to you in some part because you know what we do here at the FTC. My introductory remarks are for the benefit of the people who are here from outside the agency. So, I apologize if some of this is subject matter with which you are already familiar.

I'm going to be talking today generally about the FTC and other law enforcement's ability to combat fraud, and more specifically at the end about how we work on internet auction fraud. The remarks I am making here today are my own and do not represent the views of the Commission or any individual Commissioner.

The FTC is the federal general jurisdiction consumer protection agency. There are other federal agencies that cover more specific kinds of consumer issues, but we are a general jurisdiction agency. We have about 20 statutes and 30 regulations that we enforce.

We're a civil agency, not criminal. Sometimes people in the general public don't understand the difference, and fraud obviously can be both a civil problem and a criminal problem. We do work with criminal authorities by referring cases to them and running parallel prosecutions with them. But our objectives are to obtain redress for consumers and to stop ongoing fraud. So, we can't put anyone in jail. We bring our cases sometimes through administrative proceedings, but increasingly more civil lawsuits in federal court.

The FTC enforces a basic consumer protection statute that prohibits unfair and deceptive practices. The other statutes and regulations we have build on that same idea but give more specificity to a particular kind of problem.

Deception is pretty straightforward. It's a representation, omission or practice likely to mislead the consumer, when the consumer is acting reasonably under the circumstances, and the practice is material. This is for consumer protection only, so these laws don't really protect businesses, although occasionally a small business is treated as a consumer.

Now, unfairness is a little bit different than deception. For deception, you make a statement, it's not true and somebody gets harmed. In contrast, unfairness can be used to help us to address injuries that are being caused when there's no actual statement that is deceptive. It is some kind of practice that causes substantial injury to consumers. The harm is not outweighed by countervailing benefits, either to consumers or to competition. And the harm is not reasonably avoidable by consumers.

We use this unfairness jurisdiction to bring cases involving, for example, internet pop-ups and spyware, where people are having problems but no one is making a misrepresentation to them. You can think broadly in the auction context where this might apply to shill bidding or other things that are going on behind the scenes, where someone isn't making a misrepresentation on their listing that this baseball card is, for example, high quality when it is not. They're doing something a little bit under the surface, but we could still go after those kinds of problems, if we're able to meet the test of unfairness.

In other contexts, the FTC addresses telemarketing fraud, fraud relating to health products, and predatory lending and credit counseling schemes, among other things. It's my personal opinion that some of those are a much bigger problem than internet auction fraud, but the statistics in our Consumer Sentinel complaint database don't reflect that. Again, that's my personal opinion. Those database statistics reflect what consumers complain about, not the actual incidence of frauds.

We conduct our own investigations at the FTC. Criminal authorities tend to have investigations going on in one place and have prosecutions in another. Common investigators for the federal criminal authorities would be the FBI and the

United States Postal Inspection Service for mail fraud and wire fraud. We do our investigations in-house. We are able to use our civil subpoena power to get information from banks, shipping companies, telephone companies and ISPs. We do both covert and open investigations.

We interview consumers based on the complaints we have in our database. We do asset searches. We do undercover tapings of telemarketers. We've already talked about the Consumer Sentinel database in earlier presentations. It's available to law enforcement to search for complaints against persons committing consumer fraud.

Our main two goals at the Bureau of Consumer Protection are to stop scams and obtain redress for consumers. Sometimes when a fraud is particularly egregious, when we've been doing a covert investigation, we'll get a temporary restraining order at the onset of a case to get an asset freeze, potentially a receiver, and get repatriation of funds from off-shore. We can do that when we make a showing to the court that the fraud is egregious and that there's a risk that the funds will be removed if we don't get the temporary restraining order.

Our final relief is usually a permanent injunction against the responsible individuals, stopping them from doing the conduct in the future. Sometimes we will seek bans on certain activities if the conduct has been particularly egregious. So, you can imagine in a telemarketing case, the defendant would be prohibited from doing telemarketing at all anymore, and would not just be prohibited from giving false statements in telemarketing – telemarketing would now be off limits for him, because he has shown egregious behavior. We also usually have a requirement to disgorge profits and pay restitution to consumers, to the extent we can locate money the defendants possess.

As I said before, we often coordinate parallel investigations with criminal authorities. Sometimes we'll file our case at the same time the criminal search warrant is executed. We've been trying more in the last few years to do more work with criminal authorities. We can share our investigation materials with criminal authorities.

Now on to internet auction fraud complaints. It's definitely a problem that we have seen over many years in the Consumer Sentinel complaint database, and like previous speakers have mentioned, most of the complaints are about those sellers that just fail to ship. I think that the majority of the complaints that we end up looking at the Bureau of Consumer Protection are those who fail to ship, because those are the people who rise to the top as having multiple complaints against them. There may be complaints against sellers who have a poor quality, but the ultimate fraudster is the one who is not delivering the goods, and those are the people who receive the most complaints.

A recent trend in complaints involves scammers making "second chance offers." These are people who impersonate eBay over e-mail and sometimes by phone. They contact a losing bidder of an auction and say, "I realize you didn't win this product, but I have been authorized by eBay to offer you this product again, and we're going to do this off-site." And the losing bidders agree to this, because they wanted the good, they didn't win and they were the second highest bidder, perhaps. The person who has contacted the losing bidder will ask the losing bidder to wire money over Western Union or MoneyGram to pay for the auction item. When losing bidders do this, it is a completely unreversible transaction and they lose their money. The goods are just never shipped. We're seeing this more and more.

Usually when we have investigations into fraud on eBay, eBay has been able to provide some information to us about the people who are committing the fraud. But in cases like the second chance offer scams, eBay sometimes does not have any helpful information about these people who are working from outside the system. A common problem in finding the persons using second chance offer scams has been their use of false identities – people not providing correct contact information to their email provider, to the auction website, or to the money transmitter. It is really hard for a civil agency to find them at the end of the day.

Another interesting scam that has been used for a while is the fake escrow service that is used to give consumers a sense of comfort about where their money is being held before they get their goods, especially for high-dollar transactions. So, someone might say "I'm going to sell you this piece of artwork and you can put your money, buyer, in this great escrow service and it will be held there, and so when you get the goods, you'll tell the escrow service that they can release the money." But, lo and behold, the seller is actually controlling the "escrow service," and the buyer has a false sense of security, doesn't receive the goods, and the "escrow service" is gone and can't be contacted.

Finally, shill bidding is something with which you're all, I'm sure, familiar. A seller can inflate the price of the good he is offering by having his associates make "shill" bids for the product, forcing legitimate bidders to make higher bids to win the item in the auction.

We have a three-pronged internet auction fraud program at the Bureau of Consumer Protection. First, we do consumer education, updating consumers on trends and giving best practices advice. Second, we investigate and prosecute cases and increasingly make more referrals to local and state criminal and civil authorities. And third, we partner with industry such as eBay to obtain more information for investigations and learn more about trends, so that we can better educate consumers.

We give many kinds of education to consumers. The two main points that would compete for number one as the most important advice that we give is that consumers should protect their passwords and use safe methods of payment. Second, consumers need to learn about protections offered on auction sites and by payment providers. Third, consumers should investigate sellers, and obviously read the product description carefully and the terms of sale.

The most difficult problem that we have in pursuing auction complaints is finding the fraudulent sellers. We follow the money, and we get information from companies like eBay. We provide assistance to other law enforcement by giving them the complaints in Consumer Sentinel if they don't otherwise have access to it, giving them sample pleadings for auction fraud cases, getting consumer declarations for them and putting them in touch with each other.

That concludes my presentation. Thank you.

DR. DURBIN: All right, finally, we'll have Joe Sullivan of eBay.

PANEL PRESENTATION BY JOE SULLIVAN

~ Private Information Redacted at Speaker's Request ~

DR. DURBIN: Well, we got a bit over time here, and just in the interest of moving along, if you ordered lunch, it should be out in the hallway with your name on it. So, go ahead and acquire your food and we'll get started again very shortly.

KEYNOTE ADDRESS: THE BIGGEST AUCTION IN THE WORLD

BY PROFESSOR HAL VARIAN

Dr. Paul Pautler: I have the honor of introducing our keynote speaker today, Professor Hal Varian. Hal's the author of a large number of very successful textbooks in micro economics, as I'm sure some of you know. I actually went on Amazon the other day to find out what his intermediate micro text might cost. It turned out it was $122. For someone of my vintage who is used to textbooks that cost $30 or $40 back in the seventies, I was thinking that the $122 price tag was certainly a great reason for me to want to visit eBay and see if I could find it somewhere else.

(Laughter.)

Dr. Paul Pautler: But I'll tell you a little bit about Hal. He's the Founding Dean of the School of Information Management and Systems at UC-

Berkeley. He holds joint positions with the Department of Economics and the Business School at Berkeley. I think everybody wants him on their faculty list. He's taught at Michigan for a number of years and at Stanford and Oxford and assorted other places.

When he's not publishing economics articles or books, he writes a monthly column for the New York Times and he's covered a wide range of topics there, including a number of things that are of interest to us, gasoline pricing and gasoline taxation, which we've been spending a lot of time on recently, electricity supply and stock price bubbles. So, he covers a wide range of issues. Probably most important for us today, he's the co-author of a remarkably successful book on information economics and business strategy called Information Rules, a Strategic Guide to the Network Economy.

Please welcome our speaker, Hal Varian.

(Applause.)

~ Private Information redacted at Speaker's Request ~

(Whereupon, a brief recess was taken.)

PRESENTATIONS: ECONOMICS OF INTERNET AUCTION COMPETITION

DR. SCHMIDT: If people could take their seats, please, we'll get the session started. I'm Dave Schmidt. I'm with the Bureau of Economics here, and this first session this afternoon will be on competition of Internet auctions and our first speaker is George Deltas from the University of Illinois. So, I'll turn it over to George.

PRESENTATION: PRICING AND COMPETITION BETWEEN HETEROGENEOUS AUCTION SITES

BY PROFESSOR GEORGE DELTAS

~ Author Would Like to Note That Much of This Research is Preliminary~

PROF. DELTAS: Thanks, Dave, and thanks, Chris, and everybody else for putting together this great conference.

This is joint work with Thomas Jeitschko from Michigan State. I should say this is joint work in progress with Thomas Jeitschko because we're still very much in the middle of shaping this research and this paper.

What's the outline for this presentation? I'm going to talk very briefly about the literature and then talk about the modeling ingredients, put some algebra for the auction guys in here, and then describe pretty much all the results -- some of the results that are most important through using examples and pictures, and then end with some thoughts.

Now, what about the literature? A lot of the work in auctions focuses on seller-to-buyer interaction. This is -- you have one guy who wants to sell something to some potential bidders and then you kind of, you know, see what happens, what the equilibrium is, what things change the equilibrium. And a lot of the work that we saw presented today, this morning actually, falls into that category. You can read the papers and the books that describe much of this work.

There's a smaller literature that talks about competing sellers. This literature differs from the top one because you no longer have one seller dealing with many bidders, but there are more than one sellers who are trying to attract bidders to their product, and those sellers, you know, are also strategic players, vis-a-vis, their competition with each other.

There is a third literature, smaller but developing, and I only have one cite here, but there are more, some of the work with (inaudible) also belongs in that category, that talks about platforms or two-sided markets, and this is conceptually different than the competition between two sellers because the two-sided markets or the platforms really say there's a guy in the middle who doesn't actually -- who's not directly a party to the transaction between the sellers and the buyers, but acts as a platform where the people actually will meet. So, eBay actually is, you know, such a firm. It acts as a meeting place, as a marketplace where bidders meet with buyers. And our paper very much fits into that category.

There is a fourth literature which I think is important because it lays out some of the stylized facts that we modeled, namely the fact that even though the Internet, where you would think information should be near perfect, people should know what prices are and you would expect the law of one price to hold, there is very substantial deviation from that, that basically the Internet sellers are, in meaningful ways, differentiated with each other and, therefore, there is price dispersion on the Internet.

Now, let's talk about the modeling ingredients. So, here's the starting point. We're going to consider two sites, Site A and Site B. Sometimes we're going to consider only one site, Site A. But in general, there will be two sides, Site A and Site B. And there are a whole bunch of sellers, potential sellers for each one of those sites. And these are the guys on the left and the right here, and there are going to be M of them. Okay?

Now, what characterizes the sellers? Well, what characterizes the sellers is some cost that they have of taking the stuff that they own and placing it for sale at the site, and that cost is going to be some random draw for some distribution, G. So, the sellers are going to have different costs with the parting of their objects and selling them.

We're going to assume that the sellers that, you know, some (inaudible) will potentially sell at A and those that will potentially sell at B and later we're going to try to relax that.

Now, what about buyers? There are a whole bunch of buyers that are interested in buying stuff from those sites. Now, how many buyers are we going to have? There's going to be N times the number of objects potentially for sale in each subject. Okay? This graph is drawn for N equals two.

Now, what describes buyers? Buyers are being described by some random draw, XI, which is drawn, again, from distribution, and that XI represents tastes for browsing the Site A versus browsing the Site B, okay? And the lower the value of XI, the more likely a particular buyer, that guy with the arrow, the more likely that guy is going to browse Site A, okay, say Amazon, and the less likely it is, all things being equal, to browse Site B, say eBay.

What might generate those preferences, the sites are differentiated both by the format, okay, the actual -- what the user has to go through to buy something, but also by the features they have. So, there's the buy now feature that may be more attractive for some people than for others. Even the buy now feature itself varies between say Yahoo and eBay. It's not exactly the same feature, so people may have tastes over that. It's experiences and it's the value of buyer protection, maybe some people are more risk-adverse than other people. So, for whatever reasons, you know, potential buyers will have preferences over purchasing from the one site versus the other, all things being equal.

Now, consider this particular guy over there, that guy with the particular, draw X_I. His total cost of browsing from Site A, say from Amazon, will be CA, okay, some constant, which is the mean value -- the mean user friendliness of that site, plus theta times XI. And his cost of browsing Site B, say eBay, will be CB, which is the mean, you know, cost of CB plus theta times one minus XI. So, you can see that XI will partition in this way, but CA and CB affects the cost of all potential buyers symmetrically.

But this is not the only thing that characterizes a buyer. What characterizes a buyer is also a draw, V, from some distribution, F. And this draw is the willingness of a buyer to pay for a particular object. This is what the standard auction literature kind of has.

Now, whose object? Well, there is a guy here that's bold-faced that can potentially sell in say Amazon and he's the guy that has the object that this potential buyer wants, okay? That's where the VI comes from. There's another guy that potentially sells at B who also has that same object. So, this VI is the value that a potential buyer has to buy this guy's or this guy's object. We're going to assume that he has no value for anybody else's objects, but we can relax this at very substantial cost of complexity.

Now, who else is this potential buyer competing with? Is he competing with all the other buyers? Well, no, not every buyer is going after the exact same object. Maybe people want to buy something slightly different. So, this guy may be a direct competitor with a subset of the other potential buyers and in this case, only with one other guy who's also here, you know, bold-faced.

Now, what about the hosting site? We haven't talked about those yet. So, the hosting sites have been characterized -- their strategy is, actually, to choose a fee which it they will charge to the sellers. So, the sellers have to pay this fee to put their stuff on sale and (inaudible) is a bit more complex. You know, the fee is not necessarily a fixed fee, but for now, it's (inaudible) kind of capture the first order approximation to this, okay?

Of course, they're being characterized by the user interface, by the features they have which determine CA and determine theta, the degree of site differentiation. Those are also features for the sites they choose, but we're going to hold those fixed from the point of view of strategic introduction. Those are more (inaudible) parameters.

So, the sites have the fees. The sellers have costs of actually parting with their goods. The buyers have preferences for the two sites, all things being equal, and (inaudible) pay for the products. In equilibrium, not every seller is going to sell. Only these guys here are going to sell. The other guys are going to keep their stuff, they're going to end up having -- it's not going to be worth it to sell the stuff. And in equilibrium, some people are going to buy -- are going to try to browse Site A, these guys over here, and the others are going to browse Site B. So, that's what the market looks like.

We'll often actually look at a case of a single active site or a monopoly when there's only one site. Setting a fee will attract some of the potential buyers, some of the sellers, and the others are going to stay home and watch TV. So, this is the graphics of what the thing looks like.

Now, I'm going to put a little bit -- affix a few quantities in here in terms of -- these are the (inaudible) space that I just described. Like I said, I'm not going to go over this again except to say that, you know, Q is an important variable here. QA and QB is the fraction of the sellers that will actually put their stuff on sale at Site A and Site B. So, this is -- if you want to think of it one way, QA and QB is basically market size of, you know, Amazon or eBay or whatever names you want to give to these companies.

XC is the critical value of X when consumers have a value less than XC are going to browse Site A, those with high value of XC are going to browse Site B. That's the traditional (inaudible) type of partitioning of the market.

The way -- we have actually rigged the model in such a way that we can analyze it by fixing the number of sellers on each site to one, you know, and interpret Q as probability. You can change N to something bigger than one and think of Q as market share. It's exactly the same thing analytically.

Now, let's fix some quantities before we look at the pictures. What is the expected profit of -- we'll call this expected, it's almost expected, profit of the potential consumer who has a valuation, VI, for an object, goes to a site where there's a reserve, R, and competes against N minus one bidders? Well, this we know already from -- you know, from auction theory, it's going to be if there's one competitor with a price of -- with a valuation above the reserve and this is a second-price auction, which I forgot to say, but it's a second-price auction like it is on this side. It's going to be the highest draw out of the competing draws, and if there's nobody with a draw -- no competing bidder above the reserve, then he's going to pay just the reserve. So, this is integration over all these events.

Now, what is the entering profit of this bidder for this particular location, XI? Okay, this is going to be basically -- this is the (inaudible) with the profit of the bidder when he decides whether to go to Site A or to go to Site B. So, the only thing in -- this expression just says I'm going to take another integral, okay, this exposed profit, and I'm going to integrate over the valuation of -- the possible

valuations that I have with this item and I'm going to weigh it by the probability that we'll find N bidders where N takes
the -- competing bidders when the value takes from one to the entire market times the probability that I will find something that I want to buy. Times QA, okay?

The thicker the market is, the bigger the pay-off of going to -- the thicker the market say on eBay, the bigger the pay-off of going to eBay because the more likely that I will find at the end of the day something that I want to buy. That's why this whole thing is multiplied by QA, minus theta XI, minus EA, those are the browsing costs, and those convey from person to person because some people may value, for example, buyer protection more than others and so on and so forth.

There's going to be a similar expression from buying from Site B. So, people choose either to stay home and watch TV, browse from Site A or browse from Site B. This is what the consumers do.

Now, we're going to -- like in the traditional Hotelling case, the value of the consumer that will be just and different from buying from Site A or buying from Site B or staying home is going to be determined by the condition that says that that guy's profit, that surplus will be equal to zero. So, we're going to evaluate that profit function, and it has to be that for a given X_C, the value of somebody that has that X_C must be equal to zero. If he's going to stay home or browse or it's going to have to be equal to Website A and Website B. That's going to determine whether he shops from the one or shops from the other.

Now, potential -- so, this is -- the thing that I just covered looks like Hotelling, but the expressions are different and some things actually do look different. So, what about the sale (inaudible) problems? The seller gets revenue -- and I'm going to speed up a little bit, you know, here. The seller gets revenue from selling and the thing that you actually need to know is that the revenue is going to go up the more people that are going to be there and going to be at the site. So, the seller's value from having -- from being in sites that, you know, have more people. Why? Because if they have more people, they're going to actually end up getting more money. Of course, they also have to pay the fee.

Now, the revenue of the sellers, you know, the revenue minus the fee but not counting their entry costs, is what we're going to call the revenue function. So, the revenue function is how much money the people are going to get by selling their stuff minus the fee they pay, but we're not actually counting the cost of them actually going to the site or parting with their item.

Now, the thing that's important to know is that the revenue function depends on -- if we're going to Site A, depends on what is the critical -- what's the market size of people who go to Site A? If that's going to be determined, you can see that X is a function of both QA and QB. It's going to be determined both by how many other sellers are in that site and how many other sellers are on the competing site. Those things are going to affect the value of a particular seller selling in Site A or not selling at all.

Now, entry -- you know, the entry costs of a seller whose cost draw is at a quantile that's going to be the inverse of this cost function. This we're going to call actually the cost function. So, the cost function is an inverse of the cost distribution function.

Now, an entry equilibrium is, you know, the value of QA such that the marginal seller, the guy who is -- you know, it's the value of QA, so that under the value, the seller that has that cost realization that is an acute quanta of the cost is just indifferent between actually parting with their good or not parting with their good and keeping it themselves.

Every seller with (inaudible) QA is going to go in this market. Every seller with costs less than the inverse of this function is going to sell their stuff. Anybody with a higher cost is not going to sell their stuff. And the auction sites have a revenue which is the number of sellers that sell into that site times the fee that they charge per seller. They want to maximize that function.

So, I'm going a little bit fast here, but hopefully with some examples, we're going to make some of this (inaudible). We're going to see the analytics and we're going to get (inaudible) on this market. So, I'm going to do a first example in which the seller costs are distributed uniformly, have a uniform distribution and the buyers are going to also have uniform distributed willingness to pay for these items. So, they're both uniform.

This first graph looks at the size of the -- this is a monopoly, there's only one side here. So, this graph actually shows -- you know, plots the location of indifferent consumers, the consumer who is just willing to bid or not bid, as a function of how many sellers are participating on the website. This is -- Q is the fraction of potential sellers that are in this site. Once you see that -- and this is drawn for different parameters. What you see is that the more sellers that are on a site, the more consumers are going to be willing to pay to actually browse that site.

But also what you see is that this function is actually concave. In fact, we say in the paper that it has to be concave. It is going to be concave for our regular cases. So, it's not just (inaudible) one example, but it has to be that way.

Now, as you may -- as theta becomes smaller, as people look at the site more homogeneously, this function drops. In Hotelling as theta goes down, this function actually rises proportionately. As the cost of browsing drops uniformly, this thing also rises, but in a (inaudible) fashion.

Now, the size of this market also translates to a revenue function. The revenue function is basically how much money would a seller get by going into this -- by putting their stuff on sale if there is a particular number of other sellers also on that site, okay? So, you see that has the exact same shape. The more sellers that are in a site, the bigger the incentive for an incremental seller to also go to that site. There's a feedback effect here.

Now, to determine what's going to be the equilibrium in this market, you have to put the revenue and the costs together. So, what I've plotted here is I've plotted that same -- one of those revenue functions that shows the amount of money that a seller gets by putting the stuff on sale versus the costs of these sellers ranked again from lowest cost to highest cost. You see those two lines never intersect, which means that in this website, there can be no equilibrium in which there will be sellers and buyers interacting in that site. For no value of the number of sellers that go in this site, do the sellers make enough money to cover their costs. Now, if something happens with technology, maybe Internet becomes faster, so people now can have lower costs of browsing. Maybe they have cable instead of having dial-up modem. So, something happens that makes it easier for people to go in which shifts this revenue function up. Now, what you see is that there are three points of intersection between those two functions. One is at the origin, 00. If nobody -- if no seller shows up, then that's it, you're stuck here. The other point of intersection is here, which is unstable, which means that -- and there's a third point of intersection which is stable, you know, up there, and that's the one we're going to be considering. So, now -- so, suddenly, actually, you do have a market here.

Let me skip that. Now, what has -- so, we haven't talked about fees yet. So, what happens now is as the firm actually increases the fees, it shifts this revenue function down. It takes money from the sellers and it pockets it. EBay sort of pockets it now. But by doing so, it kind of reduces the sales of the revenue function, and now, what you see is the equilibrium, there is an intersection of the cost and the revenue function, comes at lower and lower values of Q, which means some sellers are not going to be in the market anymore. And eventually what happens, you find the optimum -- you find some situation when either removing any more sellers from this market is not worth the extra revenue you get for sellers or the traditional monopoly pay-off.

So, you might say, okay, this is what I should be doing, I should be raising my fees until, you know, at some point, I maximize my profits. Well, not so fast. I actually have shown here -- you can see the geometry of the problem, that you can actually raise your fees by just a little bit and have a catastrophic effect on your site. If you actually raise the things a little bit and you don't reverse them very quickly, what may happen is you may have no equilibrium between the revenue and the costs.

The feedback actually -- there's a negative feedback and if the sellers believe buyers are less likely to go, then sellers found it even less worth it to stay in the site and so on and so forth, the site collapses to zero.

What does the demand function look like? I'm plotting here Q versus the fee, how many sales you get versus the fee. You're not used to looking at the (inaudible) function this way. You're looking at them in this manner which is price on the vertical axis, quantity on the horizontal axis. But the demand function actually looks very weird. It's very flat and then it drops precipitously. You might say, okay, what's the big deal? Demand functions come in all shapes and forms. What I want you to do is actually compare this demand function with its counterpart if I remove the feedback effect, the fact that more sellers beget more buyers and more buyers beget more sellers.

If I remove the feedback effect, then the demand would have been given by those straight lines. Why? Because the straight lines are being plotted simply by looking at the reservation value. The seller costs are uniform and those would trace out a linear demand curve. So, I just took the linear demand curve, holding the number of sellers -- the certain number of buyers in the market fixed and you get nice straight lines. So, if you take principles, what you see is that this feedback effect, you know, makes the demand very, very elastic. That reduces the pricing power of the websites tremendously. A monopoly may not have a lot of pricing power as one might have thought likely.

Should I stop or --

UNIDENTIFIED MALE: (Inaudible).

PROF. DELTAS: So, what I want -- so, basically, this is the good news. I'm going to say the bad news. I'm going to skip a whole bunch of stuff. The bad news is that -- well, okay. I'm skipping quite a bit.

(Laughter.)

PROF. DELTAS: So, anyway, so the bad news, which I'm not going to have time to tell you, is that, you know, monopoly may not be as bad as we think because they have very little pricing power. The problem is that predation may be a much bigger problem than otherwise would be because why? Because, traditionally, we don't worry about predation, we think of one (inaudible) firm throwing out the other firm, and if you do so, then you want to raise prices to make money, then the other guy may come back in. So, why worry about predation?

Well, here, the thing that I didn't show you is that there are stable equilibria. You can push for the same fundamentals. You may have an equilibrium when sites split the market down the middle or one site has 100 percent of the market and the other has zero percent, and both of these are stable. If you manage to push one guy out of the market, then there's going to be no way that they can come back in without a huge -- paying a huge cost of marketing to get the market going again.

So, one concern may not be so big. Another concern may be actually much bigger. What about the welfare? We don't know and we don't know because there is value -- there is value in concentrating buyers and sellers in one site versus having a split. So, that makes welfare somewhat ambiguous. So, basically, in visible markets, a lot of the traditional wisdom is out the window and one has to examine pretty much any of these questions as if it was new.

Sorry for delaying. I obviously mistimed this. Thank you.

PRESENTATION DISCUSSANT -- PROFESSOR IAN GALE

PROF. GALE: Thanks. All right, this is an interesting paper. I think it's a promising paper. It demonstrates the range of possible equilibria in a particular setting with two-sided markets.

I have to say, I haven't spent a whole lot of time thinking about two-sided markets. I spent a little bit of time years ago with Department of Justice when we were thinking about credit card markets. But, in fact, two-sided markets are ubiquitous. I mean, you could have Yellow Pages, media, newspapers, credit cards, as we said, auction sites, dating services, academic journals, even, of course, academic conferences are examples of two-sided markets. Obviously, network externalities are important in all of these settings. I mean, in the market context, sellers want to have more buyers, but they want to have fewer other sellers, all else equal, and the reverse holds for buyers.

So, I mean, George did a good job of describing the model. Let me just say, when it's completed, I think this paper is going to be very useful for people who are interested in Internet auctions and, also, it's going to be useful for regulators seeking some guidance in dealing with the special issues that arise in this particular kind of two-sided market.

So, what's the main contribution of this paper? Well, obviously, on the one hand, George has shown you the technical contributions of -- these are very, very complicated models. You know, there's layer upon layer of complexity, even if you use very, very simple parameterizations, which he's done. I mean, he's -- the modeling decisions they've made have been right, just all the way down the line. I'll mention a couple things I might change, but by and large, I think they've done all the right modeling decisions.

So, the equilibrium characterization is useful, the comparative statics exercise is very useful. I think you get a lot of intuition from them.

But what's new, what's it tell us? You know, if you're a theorist and you're looking for interesting or novel results, well, you know, the results that monopoly may get higher total surplus than duopoly, you know, that's something I think we sort of understood in different contexts. We know that when there are network externalities, in fact, you can just take a simple Hoteling (phonetic) model and you can conceivably come up with a situation in which monopoly does better than duopoly if you've got a network externality component. So, you know, George pointed out that trade-off. I mean, there's less pricing competition but, of course, you've got the advantage that the bigger market confers benefits on all consumers. That's a trade-off that's understood.

You know, I'm looking forward to, when the analysis is done -- it's mostly done -- but getting some understanding of exactly the likelihood of that possibility, you know, the range of parameters and so on, just for the guidance for regulators.

Another issue that arises in two-sided markets that's different from one-sided markets is that, for example, prices above marginal cost don't need to be -- don't need to indicate market power and prices below marginal cost don't need to indicate predation. So, with credit cards, you've arguably got buyers, the consumers, paying less in marginal costs because we're being paid 1 percent back to use the credit cards. So, we're being subsidized. So, we're paying less than marginal costs and that's not necessarily seen as predatory. Although I think Judge Bork, apparently, has been on record as saying there is predation in the credit card market. But anyway, so -- but the upshot is that there's all these pricing -- the welfare results are different with two-sided markets than with one-sided markets.

Now, in this context, again, George has made this -- George and Thomas have made the sensible decision to have pricing only on one side, that is the sellers pay a listing fee, the buyers pay nothing. Now, of course, that accords, I believe, with what is actually done out in the real world. So, again, that was -- you know, that was sensible.

So, the question is, what would I do different, how would I change the modeling? Well, one thing is the possibility of multi-homing. So, the idea that in this context, buyers could participate on both markets at once. So, we know in any two-sided market, you only typically need to have multi-homing on one side. So, if everybody needs to have a Visa and a Mastercard, then merchants only need to accept one kind of card. Conversely, if merchants accept both Visa and Mastercard, you only need to have one card. So, you only need multi-homing on one side and the sensible --

(End of Tape 2, Side B)

PROF. GALE: -- attenuates that effect. So, you know, that's the one -- that's the only change that I think is really of substance.

The second one, which actually goes in the opposite direction, is just choice of location. Obviously, for good reason, you've pinned down the location on this Hoteling street for the two sites rather and, you know, that's fine. I'd be interested to know what would happen, though, if you endogenize location, and the reason is, of course, that also has an effect on this question about monopoly versus duopoly. Because, of course, when you go to monopoly, the gain is this network externality effect gets larger, but the cost is the loss of variety.

Well, of course, things are set up in such a way that we've got maximum differentiation right now. And so, actually, allowing firms to choose different locations would actually move us in the other direction and start making duopoly look better than monopoly.

That's about all I wanted to say. Again, I enjoyed reading the paper and I look forward to seeing the final version.

DR. SCHMIDT: Thanks. David Reiley from Arizona is our next speaker.

PRESENTATION: MEASURING THE BENEFITS OF SNIPING ON eBAY: EVIDENCE FROM A FIELD EXPERIMENT

BY PROFESSOR DAVID REILEY

PROF. REILEY: Thank you. So, my thanks also to the FTC and to Chris Adams, in particular, for putting all this together. It has been a really great conference so far.

I probably have been interested in online auctions longer than anybody in the room. I first got interested in 1994. EBay launched in September of '95, and so, at first, the big excitement I had in my work was trying to promote online auctions to other economists. I remember hearing, you know, as a graduate student when I'd give seminars, faculty saying, well, this will never work, people will never trust each other, there will never be any transactions like this. And I said, no, no, really, I think it is quite an exciting area to work in. And so, it's really exciting for me to see so many people doing work in this area.

Now, I see my role in life is to promote the use of field experiments. So, I'll show you an example. You've already seen one very nice one by Ginger. But I get twice as much time to promote mine as she did. So, I'll try to give you a few more details about how this kind of thing works.

First of all, I think everybody here probably knows what sniping is, but briefly, it's the strategy of bidding at the last minute. It emerges from the two important institutional features of eBay auctions. The first is proxy bidding and the second is the hard close rule as opposed to a going, going, gone rule, or I didn't even know the term before, popcorn bidding. So, Axel Ockenfels and Al Roth found prevalent sniping in eBay auctions, and when they first showed this, it was a surprise to me. In auctions typically lasting seven days, 20 percent of all last bids were in the last hour; 40 percent of computer auctions and 60 percent of antiquities auctions had a last bid in the last five minutes. So, out of 240 auctions, 89 had bids in the last minute and 29 had bids in the last 10 seconds.

The key result is that the amount of sniping varied with the rules of the auction institution. So, in Amazon auctions with the soft close rule, sniping happened much less often, and in particular, if you measured experience by feedback rating of the bidders, they tended to snipe more often on eBay when they're more experienced, but less often with more experience on Amazon, indicating that this does seem to be an equilibrium type thing that bidders are learning to do.

So, why do bidders snipe? Several theories. The one first proposed by Axel and Al was that a low revenue bidding equilibrium can be -- you know, a Nash equilibrium can be supported by the fact that last-minute bids sometimes don't get through.

The second one is that expert bidders -- actually, Pat and Ali have written about this -- expert bidders don't want to reveal their superior signals of value by bidding early.

The third one is that some naive bidders don't understand proxy bidding and sophisticated bidders take advantage of them by sniping. So, I'm willing to pay $100 for this good, but I'm pretty new to eBay and so, I, just engage in ratchet bidding, I bid $51 to get just above $50. Meanwhile, there's some sharp shooter out there like Pai, for example, who may value the good at $150 and who recognizes that there's me being naive out there bidding $51. She puts her bid of $150 at the very end of the auction and wins for $52 and, you know, I've been happy all the time thinking I was the high bidder going to win, and all of a sudden, it's gone. So, sniping could really take advantage of naive bidders.

Then, you know, some amount of late bids on eBay aren't really sniping, they're just bidders bidding on the auction that's going to end soonest and there's a lot of auctions available for the same item, so why would I bid on something that's going to end five days from now?

Dan Ariely with Axel and Al have done a lab experiment as it relates to successfully reproducing sniping in a lab setting. They found that more experienced bidders were more likely to place late bids just as in the field. The results did not depend on the probability of a late bid not getting through, which provided some evidence against their own first theory, and, you know, a quote from their paper is that "sniping may also be a best response to incremental bidding that was observed both in the field and in our experimental setting." So, that goes along with the story that I just told you about taking advantage of naive incremental bidders.

So, our research question, having seen the interesting features of timing and bid data, our research question is, can we measure benefits to sniping? You know, if the more experienced bidders on eBay are sniping more, how much benefit is there of doing it?

So, they likely are profiting from doing it, so alternatively, they might have false beliefs about it being useful or they may just think it's fun. You know, note that there is a cost to sniping. You may have to -- one way to do it is to pay for a service like we did. Esnipe.com will submit snipe bids for you automatically. You may risk to forget to bid. You know, your baby starts crying in the last 10 minutes of the auction and so you never manage to get to your computer to submit it, or, you know, even Internet congestion or inability to get online may cause the bid not to be received.

So, our field experiment involves submitting our own bids on eBay auctions to see how much better do we do when we snipe. So, this is my first time being on the buying end of the auction in a field experiment. I've done a bunch of field experiments where I was the seller and auctioned off pairs of identical goods to see what would happen with different auction institutions.

So, it turns out that some sellers frequently auction identical items at different times on the same day, and this was somewhat surprising to me, but there's actually a large number of sellers who do this. They have a big inventory and they'll list one at 9:00 a.m. and one at 4:00 p.m. and one at 7:00 p.m. and they'll use exactly the same listing for each item. So, this makes it a little bit easier to do a controlled experiment because there's less noise in the other variables in the -- you know, there's much more control actually in the variables in the auction. So, identical item description, identical seller feedback, identical auction length and so on.

The experiment treatment is doing early versus late bids. So, the early bids we submitted at least three days before the end of the auction; the last bids, we submitted in the last 10 seconds. So, 10 seconds before the auction, we asked Esnipe to submit our bid for us. So, that's the control in the treatment for each pair of items.

We chose our bid amount to be very high because we get the most information whenever we win both items. If we were to win one item and lose the other, we don't know what price we would have had to have bid in order to win. So, we picked very high bids, but, you know, we also didn't want to pick infinite levels for our own budget constraints. So, the idea is to check to see how much more cheaply we get the other one when we snipe. You might also want to measure, well, how much more often do I win the auction if I snipe, but that price basically tells you everything you might know about probability of winning, because if I get it for a lower price, then for a lower bid, I would have had a higher probability of winning. So, we think that the price in the auction, when we have a high bid, is a sufficient statistic for all the things we're interested in.

So, how do we select our auctions? I'm sorry, I forgot to acknowledge my co-author Sean Gray who worked on this with me for his undergraduate honors thesis at Arizona, and he's now a law school student at NYU.

So, we -- Sean in particular -- looked for sellers who auctioned identical items on the same day. The end times differed typically by hours, but sometimes only by seconds. We browsed by the most recently listed auctions as opposed to the auctions about to end so that we could find these things early enough to be able to submit the non-sniping bid. We looked at categories where we could estimate the resale value easily to help us set a high bid without risking huge losses. So, some categories had higher numbers of bidders per auction and some -- I saw Pat wince on this last thing, so that does come to haunt us later that there's easy resale value because that might generate less variance in prices and make it harder for us to measure things.

So, some categories had higher numbers of bidders per auction and some had lower numbers of bidders per auction. The categories include coin proof sets, DVDs, diecast cards, GameBoy games, PlayStation games, XBOX games.

So, we want to win both auctions, as I mentioned, so we use some public reference prices. EBay winning bids are generally much lower than these price lists that we're using. The values for video games and DVDs were determined by the Wal-Mart retail price. The values for the coins, we were bidding on mint condition coins, were determined by the PCGS Guide, and the values for Hot Wheels cars came from the (inaudible).com price guide. So -- and we ended up -- you know, we submitted bids that were equal to the book value, basically.

So, some data ended up being unusable. There was some removal of pairs from the data sample. There were some pairs in which we were outbid in one or both auctions. We decided to just remove these because they were noise. Now, you might worry that we'd be introducing bias if we, you know, won -- consistently won one but not the other. But it turns out that we were no more likely to have won the sniped item than the early bid item. So, if we restricted (inaudible) to this, we shouldn't get too much bias.

And then, unfortunately, we lost five more pairs of items because the seller was suspected of pirating and eBay said these were invalid auctions and they removed them from the site before we managed to grab the data. It would have been nice if we had gotten all of this data without having to pay for the 10 items, but we didn't. We weren't quick enough. So, we were left with 59 pairs of auctions whose data we could analyze and the result was that we cannot find much benefit to sniping.

Tests of difference of means between sniping prices and early bid prices, if we do it in percentage terms, the snipe auctions had prices that were 2.5 percent lower on average than the matching early bid auctions. So, there is some effect, but it's not statistically significant. I mean, at least it goes the right way, but it's not at all statistically significant and it's not very economically significant either.

In absolute terms, the sniping auctions finished about 50 cents lower on average than their matching early bid auctions, also not significant.

So, we also asked, what fraction of pairs does sniping yield a lower price than the early bidding does? Since the prices are sometimes equal across treatments, we have two different possible null hypotheses. So, if the null is that at least 50 percent of the pairs have sniping strictly favored, then we reject. If we have the null that, at most, 50 percent of the pairs had early bidding, then we don't reject that. So, not much significant benefit to sniping, even in that kind of more non-parametric test.

So, is sniping more valuable when you have fewer rival bidders? We kind of looked at it and said, well, gee, some of these auctions, you get five, six, seven bidders. Maybe that's going to get to the market price no matter what and there's less possibility for variation. So, we ranked the 59 pairs with respect to the number of bidders that actually ended up being in the auction. We divided this into two groups of 29 pairs and the low group had an average of 2.6 bidders while the high group had 6.2. No significant differences found between sniping and early bids in either of the groups.

Then we decided to ask the question, is sniping still going on just as much as it was in 2000 when Roth and Ockenfels collected their data? So, we just watched auctions of laptop computers, which was the category -- one of the categories that Al and Axel looked at in their original paper, and we found that the incidence of last-hour bids is similar, but the incidence of last-minute bids is lower.

So, in the last hour, we get 84 percent, they get 70 percent, so a bit higher. But if you look at the last five minutes, they were finding 46.7 percent of auctions having bids in the last five minutes. We only had 9 percent. And in the last 10 seconds, they had 11 percent whereas we only had 2 percent. So, it looks like, you know, a small sample, but it suggested that the incidence of sniping has gone down.

So, then I asked, well, are our categories unusual? So, we sampled bid timing data for 20 other auctions in three of our own categories, you know, in addition to laptops, and we find that we do get somewhat less last-hour bidding, but more last-minute bidding than in laptops. So, if anything, the categories that we chose to bid on had more last-minute sniping behavior than the laptops originally studied by Roth and Ockenfels. So, I'm going to be able to wrap up early and try to get us more on schedule.

Conclusions here, the previous research implied benefits to sniping. We tried to substantiate this implication through bidding on paired eBay auctions. It seemed that if there's this much more sniping by experienced guys then there ought to be a big benefit to it. We found that sniping, on average, reduced selling price by 2.54 percent, or I should say the purchase price since we're looking at it from the buyer's point of view. But it's not statistically significant. We may find statistical significance if we get a larger data set and Jeff Ely (phonetic) and Tanjim Hussain (phonetic) are currently working on a similar project.

Our analysis weakly suggests that sniping may be more useful with fewer rival bidders, but it wasn't a statistically significant difference, so I didn't bother showing you the size. And there may be less sniping than there was a few years ago. Perhaps, you know, if that's, in fact, true, perhaps it's because bidders know that the bidding pool has gotten more sophisticated over time. So, maybe the benefits to sniping have gone away because there's so many sophisticated bidders and so you can't take advantage of the naive ratchet bidders.

Thanks.

DR. SCHMIDT: Thanks. Laura Hosken from FTC will be discussing the paper.

PRESENTATION DISCUSSANT -- DR. LAURA HOSKEN

DR. HOSKEN: Thanks, Dave. So, let me preface this by saying I love field experiments. This is a really, really interesting question to me. Why do we see people bidding at the last minute when all sort of general theory says you should just bid your willingness to pay, it's a second price auction, et cetera, et cetera? So, I decided after reading David's paper and re-reading Al's paper to do my own informal survey. I called my uncle, who's retired, and sells and buys objects on eBay all the time.

So, I asked him, I said, you know, do you actually snipe your bids? You know, do you bid at the last minute? And he said, oh, yeah, all the time, I love doing that. And I was like, well, why? And I didn't even, you know -- I didn't ask any leading question, I just said, why? And his answer was -- and I quote -- "because it's fun to steal it at the last minute from somebody else."

(Laughter.)

DR. HOSKEN: So, clearly, there is, as David mentioned, possibly a little problem with who it is that is -- and even experienced bidders, my uncle's one of those who, you know, does it all the time, but he's doing it for the fun of it. That being said, Al lists a whole bunch of other strategic reasons in his paper why you might want to bid at the last minute, one of which is that it's a common values auction where there's some asymmetry in information and clearly, from the categories that David's chosen, these are not common value components to these auctions, most likely because they're posted prices. So, maybe you're not going to capture that in what you're looking for. So, that may be one reason why you're not finding that particular benefit of sniping. So, that's one of my comments.

And then my other comment is that you're performing a random experiment. My guess is that for people who are maybe trying to find collusive situations, they're purposefully choosing auctions that have fewer bidders to start with. At the very end of the auction, there may be only a few bidders left. If they know the value, there's like a large surplus that they can see available and then they choose to enter that auction. So, by doing a random experiment, which I love random experiments, you're not going to capture that either.

So, you know, that being said, what's left are people like my uncle who are doing it for the fun of it and not for the benefit of sniping. So, those are my comments on the paper, but I greatly enjoyed reading it. Thank you.

DR. SCHMIDT: Thanks. If I could ask the -- actually, I think we probably have time for a few questions. We have roving microphones for that.

UNIDENTIFIED FEMALE: One question. If you try to eliminate yourself from this by just removing, in a sense, instead of looking at the second-highest bids, looking at the third-highest bids, that would kind of keep you out of the whole story (inaudible) unless people are looking at your bid on day three and reacting to it. (Inaudible).

PROF. REILEY: I have not actually looked at that. That would definitely be interesting to do. Thank you for the suggestion.

UNIDENTIFIED MALE: Another interesting use of this data would be to -- you know, you've created some exogenous variation in the bidding behavior of one person and to trace out some stuff with the strategies. So, you see the dependent variable, you know, which is continuous and, you know, the problem we have with looking at things like auction data is everything's all simultaneous, it's all wrapped up in one, and that's just a totally different use of the data that I think would be really fun to do; that is, you submit this bid just from out of the blue and we don't normally get that sort of stuff.

UNIDENTIFIED FEMALE: I know on both of your subjects you're bidding really high value to try to win it. Is there a concern that that will encourage shill bidding? That people know that there's a certain buyer who is willing to pay extremely high on this. You may be able to tell -- like the latter of your data will be different from the beginning of your data.

PROF. REILEY: I didn't worry too much about that because we did this over, you know, so many different categories and we were only bidding in 10 to

20 auctions per category. So, we thought it was very unlikely that anybody would notice that we were the special bidder who, you know, was -- the thing is, they wouldn't be able to --

UNIDENTIFIED MALE: (Inaudible).

PROF. REILEY: Yeah, they wouldn't know the amount of my high bid. So, what you're saying is if they saw me, then they might -- the seller might actually try to bid up against me in the early bid that I've made. I considered it to be quite unlikely given that we were bidding on, you know, less than a tenth of a percent of all the auctions going on in each category. I don't -- I will check the early versus late periods of the experiment to see if anybody did pick up on this, but I'd be very surprised if they did.

DR. SCHMIDT: Could I ask the panelists for the next session to please come up and we'll get that started as soon as we can?

(Whereupon, there was a brief pause in the proceedings.)

PANEL: COMPETITION ISSUES

DR. SCHMIDT: Well, let's get started on our next panel. This is on, broadly, competition issues and we've got five panelists. So, what I'd like to do is give them each about 10 minutes to speak, and then, at the end, we can have some general questions again and then take a short break.

So, our first panelist will be Dr. Lorenzo Coppi from Charles River Associates.

PANEL PRESENTATION BY DR. LORENZO COPPI

DR. COPPI: Thank you, Dave, and thank you, also, to you, Chris, for organizing this. It's really a pleasure to be here. Let me, first of all, try and find my slides.

So, the topic of my presentation today a little bit different. It's about competition issues in B2B exchanges. The first part of the presentation will deal with the traditional analysis of competition issues in B2B exchanges, and in the second part, the chronological part, I will look at some recent development to see if we've learned anything different from the recent history of B2B exchanges.

So, what are B2B exchanges, first of all? Here's a common definition. Internet-based solution linking businesses interested in buying and selling goods or services from one another. This is a pretty big definition. It basically captures anything enabling companies to do business over the Internet. But in its vagueness, it does capture the wide variety of B2B exchanges that actually hit the market at the beginning of the 2000s.

Because of this wide variety of exchanges, the focus of the initial literature on B2B exchanges, the business literature, was on classifying and coming up with a taxonomy of B2B exchanges, and those are the six major dimensions always quoted in the literature.

Another major focus of the literature, at this time, both the business and the economic literature, was on the significant benefits for firms that B2B exchanges could generate. They could, basically, aggregate buyers and, therefore, consolidate trade. They could expand markets by creating opportunities for trading goods that wouldn't be traded otherwise. They can reduce search costs, they could reduce transaction costs, and improve the flow of information -- the vertical flow of information in the industry, thereby improving the supplies at chain management.

The main economic feature of B2B exchanges is probably that they exhibit indirect network effects. That is, the expected (inaudible) of a participant to the exchange increases with the liquidity of the exchange where liquidity captures both the number of participants to the exchange and the size in terms of transaction. And this is because a more liquid exchange lowers the expected market price patterns, therefore, a benefit to (inaudible) traders. It lowers search costs, it lowers transaction costs, and because there are significant costs of participating to the exchange and, therefore, switching costs, the B2B exchanges are commonly thought to be those markets that, because of the feedback effect typical of network effects, they may keep it to a monopoly.

Another less often commented upon feature of B2B exchanges is that they are vulnerable to coordination failure. If the final equilibrium outcome is uncertain and if it's costly to join an exchange, it pays to wait. But if everybody waits, the exchange just doesn't get started, especially if the utility of participating through the exchange is larger when the exchange is larger. But as I said, this wasn't a particular focus of the literature because I guess that coordination failure would predict that exchanges never took off. Maybe for this reason, nobody focused on that.

This is the number of exchanges in 2000, 2001, B2B exchanges really boomed. The press was hailing it the new Industrial Revolution. B2B exchanges were seen as a revolution or a completely different way of doing business. So, no coordination failure there. There was more a race to get into the market early to exploit first mover advantages and network effects. This was around the time when the FTC especially started thinking about, well, do all these B2B exchanges raise any competition issues? Do we need the new tools to review this floodgate of B2B exchanges?

Three potential issues were quickly identified. The first one is foreclosure. Exclusivity discriminatory excess may result in foreclosure. So, if buyers commit to sell exclusively through the exchange -- I'm sorry, sellers commit to sell exclusively through the exchange, these may end up foreclosing competing exchanges. If seller -- sorry, if there is discriminatory excess through the exchange, that is a seller can decide who to invite to the exchange, these can foreclose the market for other sellers. But these are often needed, also, to establish a successful exchange. Exclusivity may help solve coordination failure and also reduce any free (inaudible) -- discriminatory excess may help reduce any free (inaudible) problem that may be.

A common difficulty of assessing foreclosure in network markets is that when you receive a very high share, you don't know whether they are the result of foreclosure or the result of the natural deploy of network effects and market forces.

The other, and perhaps the most widely discussed concern was collusion. B2B exchanges increase price transparency and the low information sharing. We all know it's kind of straightforward. These increase the likelihood of collusion. But, again, we need to ask ourselves whether those are needed to get the exchange. They are linked to the benefits of the exchange. Price transparency is needed to lower search costs and information sharing is a little bit more difficult to justify, especially if it's horizontal information sharing. But if it's vertical information sharing, it's linked to the benefit of improving supply chain management.

The final point in the discussion has always been mergers and joint ventures. Here, the conclusion basically is that we don't need particular tools to think about the standard monopoly concerns of mergers and joint ventures, and also for monopsony concerns, that these are maybe a little bit more unlikely, but also fit very well in the standard analysis. I won't repeat the point that high market share may be natural in this market.

But then, all of a sudden, something changes. In two years, B2B exchanges simply busted. In 2003, out of 15 percent of the exchanges active in 2000 were still alive and all of a sudden, all research, antitrust research, economic research, business literature is simply dried up. Nobody talked about B2B exchanges anymore.

So, the state of play in a competition analysis of B2B exchanges is the one I just went through. That nothing more has been written. But can we learn something from this experience? Maybe we can. Here are the common explanations for the shake-out in the business literature. I won't review them all for time constraints, but I would like to highlight the first two.

There was very limited support from both buyers and sellers, and the second one is that the cost of joining the exchange was very significant. It involved restructuring completely the IT systems.

So, what may be the economic explanation of the shake-out? Well, first of all, as I said, the cost of joining the exchange is significant, but also, sellers are very wary of the more intense competition that B2B exchanges bring about. That benefit to buyers and, perhaps ultimately to consumers, reduce the incentives of sellers to join the exchange, and the benefits for buyers, though, are also rather uncertain because small buyers have a (inaudible) incentive to team together to get a better price, but large buyers don't really have an incentive to share their relative cost advantage with everybody else in the industry. So, they would likely not participate in the exchange.

Because sometimes the seller benefits but not the buyer, sometimes the buyers benefit but not the sellers, there is no per rate improvement and there is no mechanism to -- there are no site payments. So, there is no way to share an absolute surplus. So, basically, for this reason, probably B2B exchanges never got the acceptance that they should.

So, I will conclude with, what are the implications of this experience for the competition analysis? Well, the tools to analyze B2B exchanges are not novel. When we talk about foreclosure and collusion, it's pretty much the same analysis as everywhere else. We need to always keep in mind that the B2B exchanges have potentially very high benefits to trade off any anti-competitive effect.

But what is, I think, more important is that the B2B history has highlighted the interplay of four factors really, the cost of setting up and joining exchanges, the efficiencies in terms of marginal costs that these can bring out. But those are the pro-competitive effect. And by these, I just mean the lower margin at a given level of cost. So, sellers get less because there is a lowering of (inaudible) cost and the

anti-competitive effects.

I think that by reviewing the rich history of success and exit in this market, one can see for each industry, for each type of exchange, the interplay of these factors, and they can be put down into an equation and only if the buyer's side and the seller's side benefit from the exchange, the exchange will (inaudible). But that will tell us something about cost efficiencies, pro-competitive effects and anti-competitive effects, and ultimately, shooting for antitrust analysis of B2B exchanges.

Thank you very much.

DR. SCHMIDT: Thanks, Lorenzo. Next up is Hampton Finer from the New York AG's Office.

PANEL PRESENTATION BY DR. HAMPTON FINER

DR. FINER: I don't have a PowerPoint because of the amount of red tape it would take to actually be able to show anything to you. It's prohibitively expensive for me, especially because I'm going to talk today about at least one case that was criminal.

The bulk of our activity at the AG's Office with respect to eBay has so far focused on shill bidding cases and generally these have been referred to us by eBay and then eBay will produce typically all of the data that underlies those matters, along with some of their understanding or expectations about who might have been a shill bidder, and then we conduct our own investigation and find out all kinds of interesting things, ultimately take that data and try to come up with some estimates of what the harm to the other bidders in the auctions were and have some sort of civil action, or at least, in one of these cases, have a criminal action and get the money back and give it to the aggrieved parties.

So, we've had kind of both kinds of shilling. So, I mean, you know, these are -- the two kinds that are out there, and most people in the room are probably familiar with that, are shilling that's used to sort of authenticate fraudulent or counterfeit items. So, we had one case that involved counterfeit art. They were very, very bad counterfeits. Apparently, they were being made by a high school student and going to the woods and, you know, it was just --

(Laughter.)

DR. FINER: Of course, it was a criminal matter ultimately, so we had to send this to all the experts in the world who were really happy. One person could tell from the eBay ad, which was like this really blurry picture, that it was counterfeit just like off the top of his head. It was amazing. He was like in Denmark and he just knew immediately that it was counterfeit. But further investigation found out this person also had live auctions which he would shill, and the way he would shill them is he'd also act as the auctioneer and he would just throw in another bid and just raise it.

This was perfectly typical behavior to him. He didn't seem to have much compunction. He was also a habitual and continual counterfeiter. He had been doing this for years and had actually served a significant amount of prison time earlier and just can't seem to stop himself from counterfeiting art.

It had all the classic kind of things you'd expect. So, if you're trying to authenticate something, it doesn't do you much good to snipe, obviously, because no one will see that bid. So, he would have a confederate that would be like an employee or himself using an employee's user ID or himself using one of his other IDs. Again, eBay very easy to snipe because you can have -- I don't want to accuse eBay of anything, but you can have a lot of different identities. So, it's very difficult for someone to detect sniping unless they have a lot of your data.

So, in this instance, most of these shill bids were usually the first bid. So, he was putting a bid of some reasonable size on a counterfeit item and that was obviously meant to convey the expectation that this was actually real, and there have been some higher profile cases along these lines.

In any event, the damages in that case were very easy because, as far as we're concerned, they were worth zero. So, anything that anybody paid, he had to give back. So, we didn't have to worry about any kind of interplay or common values or whether he was doing the second type of shill bidding, which involved sports cards. There was a guy who actually won the lottery, like a PowerBall or a New

York State Lottery, and quit his job, bought himself a PT Cruiser and bought a franchise to sell a particular branded kind of sports memorabilia by a company called Steiner. They have a lot of exclusive agreements with New York Yankees. And he was out on Long Island, and so, he bought this franchise with his lottery winnings and started to open up kind of an eBay shop. As far as we knew, he had sort of the exclusive eBay shop for Steiner memorabilia.

Now, this doesn't mean a huge amount to me, but this is particularly desirable memorabilia. I think Derek Jeter had an exclusive with him, so there was a lot of stuff like that. Anyway, he has his family bid on stuff generally at the end. So, this was either kind of the reserve updating kind of argument of, you know, they'd hit your stated reserve and then you'd say, okay, well, let's see if I can get them to go a little bit higher. Unfortunately, that wasn't the case as much as we might have hoped.

Usually what he was doing was just taking it off the market. I mean, he didn't want to see it go -- it was a no reserve auction. He didn't want to see it go for less than he was paying for it, and so, he would just step in or have his daughter or his son-in-law or someone like that step in at the last moment and win the auction. So, that happened a fair amount. Although sometimes he was successful in inducing, I suppose, others to raise their bids.

Now, you know, I've read some of the theoretical literature on this and I understand that -- you know, by some interpretations, this isn't necessarily going to lead to harm, at least over several of these, but, you know, nonetheless, our damages analysis was, you know, if the bid actually impacted the selling price, we would sort of pretend that it never happened and then, you know, basically calculate what the price would have been, but for his activities, if it was sufficiently late in the auction. So, if they were showing kind of early in the auction and ultimately another non-confederate came in, you know, we would sort of ignore the shilling and not worry about the bidding history having anything to do with the final selling price. But that was that.

So, we ended up with about 40 auctions in which it actually had an impact. It was $1,200. It wasn't very much money. We wanted to take the PT Cruiser, but that wasn't going to happen. In the first instance, he had to take a criminal plea, this was civil, and it was a very colorful family for a variety of reasons. We had them in and we talked to them and, you know, they kind of all hated each other and I didn't know that my dad was using my user ID to put in -- you know, it was that kind of thing.

(Laughter.)

DR. FINER: But, ultimately, it was an interesting case. So, I guess the question that's probably in all of your minds is, well, what does shill bidding have to do with competition? And I guess the answer is, I don't know exactly, but I can tell you that I'm assured that under New York State law, creating the illusion of competition is, in fact, an antitrust violation. Under our Donnelly Act, it is actually illegal to do that. But that's not necessarily an orthodox antitrust analysis, at least not the way I was taught. I'm actually an alumnus of the FTC and it's not something I would think of first, but apparently it is of some concern in New York State.

So, there's sort of your anti -- that's where it gets in the anti-trust law. But mostly this is really just a fraud and I think a relatively garden variety fraud. A more interesting question is, why does eBay care, and there's been several papers on this. Obviously, shill bidding tends to move the prices up and eBay gets a fee that's based on price, then one might imagine that they would tacitly encourage shill bidding. At the same time, you don't want to discourage people from coming in and thinking that it's fair. I'm sure the theorists in the audience have all kinds of explanations and have actually done all kinds of good work in this area.

So, I think rest assured that it's illegal under New York State law and we will continue to pursue this and are continuing to pursue these kinds of cases. In particular, the big puzzle for antitrust is that often this doesn't involve any kind of conspiracy, certainly not among competitors, but not even among two people. If I'm conspiring with myself, am I really conspiring to do anything? Antitrust typically requires at least two entities, and these are interesting cases because they only have one. Again, the illusion of competition question, insofar as that's an antitrust violation, would come into play regardless of the number of players here.

Finally, I had the question of how common is shill bidding? So, not being as familiar with all this great literature and realizing how rusty my econometrics skills are now, seven years out of graduate school, living as the single economist amongst thousands of attorneys, I just decided to be really simple and thought I was being really clever until Patrick talked about iTunes, and I looked at gift cards and I said, well -- there's a wholly different reason. There's an article in (inaudible) which I found quite interesting. Why does anybody want a gift card? It was completely unrelated.

But I wanted to know how much do people really value gift cards? So, I wanted to see what percentage they went for based on face and it turns out it's 88 percent, if anybody's wondering, and it's 99 percent for American Express gift cheques, which is essentially cash. So, there's your probability of malfeasance or something like that, which I thought was quite remarkable in its own right. Gas is 98 -- gas gift checks are 98 percent of face value --

(End of Tape 3, Side A)

DR. FINER: -- 2,000 auctions. Lo and behold, between 4 and 5 percent of them were 100 percent or more than face value. Well, I mean, there's my lower bound on how much shilling there is. I went and investigated them and there were clear repeat listings. So, that's actually pretty high, kind of (inaudible) high actually. But I'm sure there's all kinds of problems with that analysis and I'd love to hear better ways of doing that, but this was just my lower bound. Obviously, shill bids that are used for different purposes, not just to take it off the market or -- you know, I think that would be largely what this was was an effort to take it off the market, but it could have been sort of a mistaken (inaudible).

But successful shills where they manage to induce someone to go up, obviously, wouldn't be captured by that and the winning bid might come in a little bit lower somehow. But I was quite amazed. Incidentally, these are all free shipping auctions and none of them were collectibles. So, someone's probably wondering if I got any Disney or something like that. But it really wasn't that bad.

I think we're working on other things. We do bidding ring cases a fair amount. I see my time is running out. But we've done some bidding rings. We did a stamp bidding ring, which wasn't on eBay, but it might as well have been, it was actually at live auctions and it was full of colorful characters, too. I really do recommend, if anybody's in the law enforcement world who can do this kind of thing -- I know the FTC generally can't -- it's quite interesting and there's lots of (inaudible) data and eBay is super-cooperative and super-interested in sharing this kind of data in general, and so, we get all the kind of gory details, including the addresses and the locations and the names and all the kind of secrets about the bidders and the sellers, which is quite rich.

So, you know, I imagine that there will be several cases of this type in the future, and thanks for listening to my story.

DR. SCHMIDT: Next up is Lawrence Coffin from Beantown Trading Post.

PANEL PRESENTATION BY LAWRENCE COFFIN

MR. COFFIN: I'm talking to you today more as a representative of the eBay seller side of things, as opposed to someone who's actually out here studying eBay. We're a trading assistant business up in Boston, and one of the things that I have been very interested in is the regulation of eBay sellers. In the last couple years, there seems to be an increase in states trying to step in and regulate eBay selling. In particular, in the trading assistant business, where it's very easily identifiable who the sellers are, they're starting to step in and apply different regulations on eBay sellers.

There are two types of regulations that are being applied. The first is for auctioneers, what are typically applied for auctioneers. They're requiring that eBay sellers essentially become licensed as auctioneers. The other type of regulation is second-hand dealer laws, which typically are applied to pawn shops.

In the case of auctioneers, it's somewhat complicated as far as trying to get licensed as an auctioneer. I mean, you have to take courses, you have to become an auctioneer, which costs anywhere from $1,000 to $2,000 and takes about eight to ten days as a course. You have to get licensed in the state or the city that you're in that's saying that you're trained and everything. Some areas require that you put up a bond, some areas require that you do all your transactions through an escrow account. Most states also require that you have a continuing education. This means that every two to three years, you have to go in and take a couple days worth of courses and everything.

For second-hand dealers, it's not so much an upfront cost. There isn't any kind of training or anything involved, but there's an fairly high ongoing cost or what seems like a fairly high cost. First of all, you have to record detailed information about the people who are bringing you things to sell, such as where they live, driver's license number, stuff like that. Some states even go so far as requiring thumb prints, fingerprints of your people. You also have to report daily to the police everything that you bring in. You keep a log of everything you bring in, send it off to the police. And most locations require that you hold the items for 15 or 30 days before you're going to sell it.

So, not a huge upfront cost, but these are things that we see tend to really discourage people from using our services. Usually when they come in, they want to sell something fairly quickly. Already it's taking two or three weeks for us to actually get the item listed, sold, get the money from it, make sure the buyers are okay with it, and then pay the seller. So, if we had to put it on another 15 or 30 days, a lot of people are going to walk away from that.

A lot of people also find that the fingerprinting is very much an invasion of privacy. Also, you have to look at who this is impacting and who these regulations are applying to. There are probably 15,000 registered trading assistants, but they span the gamut from your franchise stores that have locations in every city and every state down to someone who sells maybe 10 things a month and they are offering their trading services as an aside thing.

Some of these people are working out of their houses and things like that. Even small businesses, people that are running this as a business on a serious basis start out in their homes. You know, they clear out all their personal stuff, fill it full of boxes and shipping stuff and shelving and they work from there. It's already awkward enough when people come in for you to say, well, can I have your driver's license number. If you have to pull out an inkpad and say, could you please give me your thumb print, standing in the middle of all your boxes and everything, people are kind of a bit hesitant about that.

So, primarily, we’re seeing these regulations being applied to drop-off stores first because they’re the most easy to identify, the most easy to target. They have a fixed place and signs outside that say what they do. Above that, then we also see just general trading assistants.. It’s really not difficult to find them. They generally advertise. EBay has a directory listing where you can go in, type in your zip code, see everybody who’s registered as a trading assistant within 10 miles of you.

Some of the regulations are actually written to be very broad and they have the potential of being applied to essentially all eBay sellers, not just people who take things on consignment. Like the second-hand dealer laws also tend to apply to people who buy stuff at yard sales and things like that. So, these regulations do have a lot of potential of impacting a lot of eBay sellers.

We’ve seen maybe two or three states in the last year who really were starting to look at it. Right now, there’s about five states and one town that have active laws that have actually said, yes, eBay sellers have to follow these laws. There are two states that are looking into it. There’s one state where it was actually defeated. In Florida, they tried to apply these regulations. They took it to court and the judge said, no, it doesn’t apply to this person. There’s another state where that happened; where they were trying to expand the auctioneer laws to say that everyone on eBay has to sell as an auctioneer. That got turned around and got defeated.

But we're seeing that one of the states is expanding its second-hand dealer laws to include consignment shops. While it's not clear whether they're going to target eBay sellers, it definitely opens up the door to that some more.

So, anyway, this is something that we feel, as sellers, is starting to affect our ability to compete and our ability to continue our businesses as they are. So, I just wanted to present that to all of you.

DR. SCHMIDT: Thank you. Next up is Bob Marshall from Penn State and Bates White.

PANEL PRESENTATION BY DR. ROBERT C. MARSHALL

DR. MARSHALL: Thanks to Chris Adams and the FTC for inviting me. I'm not on the right slide here, so -- am I going the wrong way? Okay, sorry. For my co-authors in the audience, they're used to this technological incompetence, but for others, this should be surprising.

(Laughter.)

DR. MARSHALL: There, that's it. It's rubbing off. Thank you.

So, I'm at Penn State, but I'm in town for a couple of years at a firm called Bates White. Five years ago, six years ago, that was a firm of five people and now it's 130 and two of the founding partners are my former Ph.D. students. So, they asked me to come to town for a couple years.

I wanted to talk about an issue that I've thought about and written about through much of my career on bidder collusion. I didn't take the thoughts here to be -- or the title of the conference to be specific to just eBay kind of auctions. So, this briefcase I walked up here with, I bought this on eBay. I've been very pleased with this briefcase. I don't think I had any concerns that there were colluding bidders when I went out to buy this briefcase and I don't think the seller should have been concerned about that and I wasn't approached by anybody to collude and I'm very pleased with the purchase of this briefcase.

But I would -- there are many other environments to think about. The Federal Government is engaged in conducting lots of auctions and procurements and moving ever increasingly to Internet platforms. As Lorenzo was talking about a moment ago, there are many B2B auction procurement formats. And for me to take an experience that I had from eBay and say, if I was asked to help a Federal agency with the design of an auction and say, well, I don't think collusion is really a problem to be worried about, I would have a great cautionary word about that.

So, I think it's -- my opinion -- one of the biggest threats to revenue or acquisition at lower cost that exists out there. So, let's think about the theoretical literature for a minute. We think about risk aversion and bidder value affiliation, we're playing around with the first and second order statistic here. That's what it is. When you're talking about collusion, you're not talking about the first or second order statistic. You're talking about like the seventh or eighth. I mean, there's some significant problems there if you have not designed appropriately to take account of the potential for collusion.

Now, you may say, well, a lot of cartels out there, they're about other things than bid rigging. Market share allocations, break up the world into various components, this firm gets this percent, another firm gets another percent. Geographic allocations, customer allocations. Big international cartels, they work that way. At the end of the day, they've got to supply, as vendors, a particular customer and that customer is going to run a competitive procurement and they're going to be asking for bids and those firms that are colluding, whether it's through a market share agreement or a geographic allocation or a customer allocation, have got to rig bids. There's no other way to do it. So, I think it's a very big threat and to not pay attention to it in the design is potentially a mistake.

So, what's the specifics about Internet auctions? I haven't said anything so far that's different about any other auction format. Why pay attention to this issue of collusion with Internet auctions? Well -- did I skip one already?

There's a sequence of questions here that we might think about. Has the Internet created opportunities for collusion that did not exist before? So, keep in mind, any time you've created a new rule or you've maybe provided a new piece of information, okay, you may say, well, there's a revenue or cost reason for doing that, but I'll come back in a moment as to why that's something to be very careful about when you're focused on the issue of collusion.

Has the Internet inhibited collusion that used to exist? How should Internet auction procurement design be impacted by the potential for collusion? What information can be retained when using the Internet for auction procurements? Will knowledge that such information is being retained impact collusion? So, he's all crazy about collusion. Well, these are questions that I think hard about and I think are worth thinking hard about.

You go to the European Commission and read their decisions of the last decade, which is wonderful reading. I'd rather read that stuff than most novels that I pick up. These are just -- the DOJ and the FTC don't write things like this and it's just features of perhaps U.S. law and the way it works that prevents that kind of details from being provided. But these are wonderfully rich descriptions of how cartels actually function and what they're doing and how they're thinking. You read that and you say, well, what industry in the last 10 years isn't this impacting? So, I just think these are important questions to consider.

So, Lorenzo touched on this point, but I would like to say it very forcefully, colluding bidders like transparency. That's a good thing if you're colluding. So, let's think about how you might collude with me in suppressing our bids and obtaining a better pay-off at an auction and what that means potentially for how we have to interact with one another. Well, I'd like to do something like the following perhaps, I'll submit a bid and you suppress your bid completely. Okay? Well, I would like to know that you suppressed your bid completely. I'd like to observe the suppression of your bid completely. So, you start putting out -- we're going to put out all the bids out there for transparency reasons. So, you'll hear a lot of this in Federal procurement.

For transparency reasons, put all the bids out there. We can put more information out there because it's the Internet now. We've got the time at which the bid was submitted, every single bid that was submitted at every single time. We can put it all out there. Transparency. Well, that's just one piece of information after another that colluding bidders can condition upon to enforce their behaviors. This is a concern in my opinion.

So, I would say that this need for transparency for fairness and openness, you'll get a lot of support from the vendor community potentially for that. But I'd, again, be very careful about that.

Inter-bidder communication. Once you start talking about making it fully open as to the bids and the sequence of the bids and what those bids are, we've already seen, in some auction contexts, how those bids get used to communicate. The AB blocks of the FCC auctions was a good example of that.

But you don't have to put all this information out there. So, why is the Federal Government, for example, when it conducts procurements, why is it so concerned about making things fair and open? Well, there's an issue about what is the bid taker doing here? How are these bids evaluated? Roll all the cards face up so that we can see what the bid taker was doing. If you can instead put this into some kind of machinery through the Internet and say, okay, you know, you don't get to see anything because all the bids are going into a piece of machinery, all the information's going into a piece of machinery, what's going to come out is you won or lost, and if you won, here's what you are going to get paid if you provide a service or here's what you're going to pay to obtain the commodity. That's all that, you know, potentially needs to be done.

And you say, well, as a bidder here, I'm concerned about their bid taker (inaudible). You say, no, no, no, we have an automated mechanism now.

Now, what is the other thing this allows you to do? Because this is a huge problem. You go talk to a bidder and they have this allegation of collusion by the vendors who've been providing them. You go, okay, let me have all your data. They're going to say, what do you mean by that question? Well, all the data you have from all the procurements you've conducted. Well, we know the winner and we have some recollection of what happened there. We don't know how many people bid. Well, this is a remarkable mechanism for keeping track of everything now. Everything can be retained. You can provide very little information to the bidders and retain everything they told you.

Now, why is that relevant in the long run? Well, if you let bidders know you're going to do that, that seems to me to be a potential deterrent to any behavior they might engage in in the future, or currently with regard to collusion. So, I would say that the conditioning on non-collusive behavior and design is naive and I would really encourage that to be something that's at the -- if you know you don't have to worry about it, fine, then make that determination first and move on. But if -- to weigh that question first, I think, is a very important one.

And then the last bullet here is to emphasize the use of schemes that require that the ring -- that's a cartel -- to change the bids of all bidders, especially the one with the highest value or the lowest cost. The schemes which allow that bidder with say an auction with highest value to do just what they were doing, acting non-cooperatively, are well-known to be very susceptible to potential collusion.

Now, this is unashamed marketing. So, at Penn State, we have started a center for the study of auctions, procurements and competition policy. So, that Center is launched. There's the website. So, these faculty include Vijay Krishna (phonetic), got that wonderful book on auction theory, and Quang Vuong and Isabelle Perrigne, who have been at Penn State for three years, although people look surprised still when I say that.

(Laughter.)

DR. MARSHALL: And Mark Roberts and Jim Tybout and Jim Jordan. So, this is a remarkable group of faculty and we're trying to build bridges to the private and public sectors with this center and look forward to thinking about very important issues for Federal procurement and a lot of other sectors in the economy.

Let me just say the following. I was at the Acquisition Advisory Panel this morning of OMB making a statement to them. I was asked to come and make a presentation about Federal procurement. Their questions were really about just tell us why competition is a good thing. Now, you know, this is -- again, I come back to, if that kind of concept needs to be explained very clearly, you really have to, in my opinion, say let's think hard about the advice we offer as economists with regard to design of auctions and procurements and do we really need to pay attention to collusion or not because we're not -- the people who get to design these things are looking to us for that kind of input. It's not just intuitive for them.

Thank you.

DR. SCHMIDT: Okay, and our final speaker for this session is Kulpreet Rana from Google.

PANEL PRESENTATION BY KULPREET RANA

~ Private Information Redacted at Speaker's Request ~

DR. SCHMIDT: Thanks. We might have time for a question or two if anybody has something of general interest. Otherwise, I'd encourage you to maybe talk to the panelists in the upcoming break. Are there any questions?

(No response.)

DR. SCHMIDT: Okay. Why don't we take five or ten minutes and start the next session? Thanks to all the panelists.

(Applause.)

(End of Tape 3, Side B)

PRESENTATIONS: INFERENCE FROM BID DATA

DR. ADAMS: My name is Chris Adams. Everybody's been thanking me all day, but I want to thank everybody for coming today and being so great in answering my emails and getting back to me and reading my continuous stream of emails the last week or so.

What we're going to get going on is the last session or the last group session, which is probably the stuff which is most interesting to me, talking about bid data and how we can use it both in research and practice. We're going to start with Robert Zeithammer.

PRESENTATION: FORWARD-LOOKING BIDDING IN ONLINE AUCTIONS

BY DR. ROBERT ZEITHAMMER

DR. ZEITHAMMER: Thank you, Chris. I'm going to talk today about my paper that's mostly focused on demand. In particular, I'm focused on demand in very specific eBay categories that exhibit private values and moreover, that tend to have unit demand on the buyer's side.

I'm going to try to persuade you that the demand side of eBay in these kind of categories is actually very dynamic, and not only should it be dynamic, but it also seems to be dynamic when we actually look at the behaviors. An example of such a category is consumer durables, and more specifically, the digital camera. Suppose I go to buy a digital camera on eBay. I search digital camera, maybe point-and-shoot digital camera, what have you, and out comes this particular scenario. Sure enough, there's a lot of digital cameras to choose from.

To go further, lets say I like the Canon S-30. There is an auction ending in 15 minutes. The little box on the left is a stylized webpage that I'm looking at, and there is another Canon S-30 ending in 47 minutes and so on. Similar situations could be encountered in things like electronics, movies, or computers. I notice that the unit demand assumption makes sense. As a private buyer, I'm looking to buy one unit.

Now, not everybody likes Canon S-30s, some people maybe like the Canon S-40, some people might like Olympus, etcetera. Notice that while all these auctions are listed simultaneously, they're not truly simultaneous. In fact, they're sequential. They're ordered by their ending time. And in this particular example, there are really interlaced sequences of auctions for essentially identical units, barring used versus new and things like that.

So, I'm going to try to motivate the following conceptualization of eBay, a sequence of auctions with overlapping information. Now, when I say sequence, I mean that the auctions that I'm going to be thinking about will be few bid auctions, happening at the very end, much like the talk we've already heard today about sniping.

So, in fact, auction one gets started at this point in time, and I don't really know what happens in it, but technically or theoretically, nothing should happen until time T1, at which the people put their bids in. Notice that at time T1, auction two is already known and two questions arise. One, how should the bidders bid? How should they incorporate sort of forward-seeing information? And second, do the eBay bidders seem to oblige and behave accordingly?

Now, I'm going to go through an exercise of developing the model of bidding behavior. The model is what I would consider rigorous, even though you could probably guess at this point what the outcome's going to be. Surely, the better the future for me, the less I'll bid today. Well, that intuition is coming from some sort of a single agent dynamic programming idea, and what I'm going to try to persuade you, is that it sometimes holds in situations where the bidders are actually competing against each other in some sort of a symmetric equilibrium. They're all actually trying to do the same things. They're all trying to bid less and it all sort of works out.

Another thing I'd like to point out, and actually relate to the paper by David Reiley, is that this is not that last-second type of sniping. It's more towards the end kind of type. They don't seem to explain it very much, the reasons for these naive bidders to bid towards the end, because if you just bid in the last hour, there's plenty of time for them to outbid you.

However, something like bidding within the last hour, not necessarily the last second, would make sense in this situation, right? Because if this forward-seeing information is valuable, you might as well wait towards the end when you have the correct information to make the right bid. So, there's another explanation for sniping that's sort of roughly consistent inside this model.

But let's first step to the model. Before I go on, I'd like to point out two sorts of future auctions that are very different from my point of view. I'm looking at the first auction for my camera, Canon S-30, auction one. Already, I see auction two, and I'm going to call that the forward-seeing auction. So, in some sense, I'm playing with the titles like forward-seeing behavior as opposed to forward-looking behavior. And then, of course, eBay's not going to end when auction two ends. There's going to be some future auction that's not even listed yet. But surely there's going to be another camera sold in the future, and that's the sort of potential, yet unseen, auction described in the recent paper by Jofre-Bonet and Pesendorfer.

How does that relate to the literature? Well, I'm going to be, in some sense, extending the model of Milgrom and Weber, but not the one you're thinking of. It's part two, which was written in '82 and actually finally published in '99, which deals with finite sequences of identical units. Now, of course, in finite sequences of identical units, there is no use for future information because you know it's all identical. The finite model is actually very nice. It helps them to not worry about replenishment of the bidder pool, and they just have results based on order statistics. It very nicely shows what a symmetric bidding equilibrium looks like.

Another stream, which is perhaps more applicable to eBay, is the idea of stochastically equivalent units. The idea is that the units are identical only in expectation, but there is some variation, and these are the yet unseen ones. You don't know what's next. So, in the camera example, I'm bidding on my Canon S-30. I know it's a Canon S-30, and I'm told that there's going to be an auction tomorrow for another digital camera. Which one, I don't know. So, that's the kind of situation.

In the finite sequence, Jofre-Bonet and Pesendorfer look at an infinite horizon version of this model, roughly. I'm sure there are some other differences between Jofre-Bonet and (inaudible), but no information about future auctions, and that has two implications. First of all, it's interesting from this overlapping point of view. Second, notice that since these units are identical and everybody has the same information, this is a symmetric and independent future. So, in some sense, you can use backward induction to solve this game. This is nice and it's going to be something that's going to disappear when I go to my model.

Finally, Ian Gale was here a while ago, I don't see him anymore, but he is probably on the forefront of looking at situations where not all the units are the same and there is useful information about the future. They were only able to solve two auctions and two bidders because that's already pretty tricky. .

Now, the tricky thing is, if V2 is much bigger than V1, there's some issues of disposal. Maybe I would like to just throw away that first unit and bid again, even though I paid for it. But if I really like V2 much more and I happen to win V1, even for free, I might want to throw it away and that's tricky. This is, of course, very hard to extend to many auctions and it's actually fairly hard to extend to multiple bidders beyond two. Multiple auctions is really where it all breaks down and various asymmetries emerged.

In contrast, I'm going to only allow the values of V1 and V2 to be either zero or V. So, I'm going to look at this as a finite set of product types. I guess I can move on to my model.

So, there's going to be an infinite sequence of second price auctions. Why second price auctions? Because the benchmark is easy to see and it's somewhat easy to solve. There will be, at existing points in continuous time with some times in between that would be varying, and each of these auctions is going to sell only a unit of some discrete number of types. Here, I'm looking at a two-type example because that's enough to get all the information across. I'm going to pool substantially across types, and set no reserve just because that's easier.

The idea is that the bidders only want one of these types. So for me, the type is the Canon S-30. For somebody else, his type may be Olympus D-40, and the value for other cameras is just zero. That is obviously a strong assumption, but it helps me out in actually showing that the model has a solution. Disposal issues don't conflict or make everything too complicated.

So, I have this unit demand for only my desired type. I'm going to assume independent private value on that, but at one unit, let that be continuous, that's not a problem. In some sense, the future information is going to be just information about the desirability of the next unit. I will know whether or not I want the next unit, and second, how long it takes to get to the next unit. It's a little bit like knowing if the next auction in the sequence is Canon S-30 and if it is an hour from now or a day from now or what.

Everyone (inaudible) because of the infinite rising model and, again, to maintain a symmetric equilibrium, I'm going to have to assume there's no memory. All sorts of asymmetries emerge when people can condition their bids on the outcomes of yesterday's auctions. This happens because yesterday's auctions are actually informative about tomorrow's competition. That problem has been nicely outlined already by Milgrom and Weber.

So, these are interlaced sequences, and in this case, two identical good auctions with non-overlapping populations. Could it be made overlapping? Yes. Is it easier to do it with non-overlapping? Of course. Some bidder replacement is essential here. Notice that if I only replace the winner, we would end up with a steady group of N bidders with values zero. Then the next replacement would come in, win, and leave, and this process would continue forever. This is something that I learned the hard way, because I was trying to simulate it and it just kept converging to zero. So, there's some necessity to exit and to continuously kick people out of the market and replenish them with new bidders. You cannot just skim off the top.

Finally the innovation here is forward-seeing information. This is a picture of that situation. You see only one period ahead. These are the desirability indicators, and I'd just like to focus on the bidding strategy here. The bidding strategy solves, as a function of do I desire a current one, tomorrow's one, how long does it take, and my private value. Of course, the surplus that I'm going to get if I lose today depends on whom I'm going to lose to. The higher that person is, the higher I think the surplus is going to be. I have to think about it because the bidder pool evolves. Basically, I'm looking at literally discrete N people, and whoever loses today is still going to be around tomorrow. So, if today's high bidder is very high, then that means that the upper bound on the remaining bidder is higher and then tomorrow, the bidders are likely to be higher. So, I have to account for that.

The good news about this first condition is that what I'm maximizing over is only the limits of the integration. This is with second price (inaudible) auction coming in and a truth revealing mechanism. In some sense, my bid here, only affects my probability of winning or losing. So, same intuition comes through, the pivotal intuition.

First, the conditions tell me to bid -- well, first of all, zero because this is not my desired item. This is pretty important. If it is my desired item, then my bid should be my valuation minus the expected surplus I would get if I lost to somebody exactly like me. So, that's conditioning on losing to somebody just like me, and I call that pivotal thinking. You have to bid as if you're about to lose in a tie to a bidder just like you. This actually comes through in the original Milgrom and Weber Model as well, but this is how it manifests here.

Of course, another thing you notice is that the surplus is always less than V, and so, you bid somewhere between zero and V. There is bid shading, the idea that the higher the supply, the lower the price that comes out in the end.

It's important to notice in these kinds of markets, how does the supply affect price? Well, through the forward-looking strategies of the bidders. If the bidders are naive and are bidding the valuations, this market would not work very well. It would not adjust appropriately for the amount of goods being sold. In this market, the burden of proper pricing is placed on the bidders.

Of course, I have only told you some sort of best response. In addition, I have to then show that there is such a surplus function that satisfies the development equation. That is the correct assessment of the surplus condition on losing to somebody just like me. It becomes a whole big mess. You cannot solve this explicitly for even the simplest distribution you can think of. However, you can show using fixed point theorems and function space that this auction exists and this is well-behaved. You can also play in Matlab and estimate that function if you wish. But my question to you would be, could this be a basis for structural approach? I'll leave all structural discussion right there and move on to my approach, which is very much the opposite.

What I'm going to take to the data are properties of the equilibrium bidding. I can show that in the best response situation, through analysis of the surplus function, it must be true that bids are positive only when I desire the type. This is not surprising. I'm just making sure that all these intuitions from best response carries through to equilibrium.

My empirical strategy is going to be to assume the first line. That's how I'm going to identify what's your type. If I see you bidding a positive amount on something, that is if I see you participating in some auction, you must like that type. Since this is public information, it's really not necessary for me to do the usual structural thing and try to deduce what you're underlying valuation is. In some sense, I'm not going to integrate it out, but I'm going to average over the V. In particular, I'm going to look at orders that sticks in V, which are actually observed.

I have been talking about one period look ahead and only two types. There are more than that. I'm going to look at many types and several periods look ahead. I'm going to look at five-period look ahead in my actual model. The important difference between the timing and the types is that timing is type independent, whereas I care more or less about different types than somebody else.

So this public information means different things to different bidders about types, but the timing information means the same to everybody.

I'm going to basically hold certain things about this state constant (inaudible), that's going to give me some sort of on average prediction. If something is true for every valuation, V, then it also must be true for all the statistics of the valuations within each auction, keeping the number of bidders constant. This is going to be important. Then I note that in eBay data, we actually observe the second-highest bid, the price, and eBay, itself, observes the highest bid.

All the others, they may be truncated in some funny way because the person coming in is too late to the party and the minimum bid has already risen above their value, and we never see them. But the first person and the second person we always see. Since people cannot outbid themselves, these are different people we're talking about.

So, what I'm going to do is run some simple regressions, putting in some variables. Type specific, I'm going to look at time until next auction of the same type. I'm going to look at indicator function of whether or not the current type we're bidding on now is offered in the next five auctions, and then I'm going to actually do this forward (inaudible). Is it offered only at one auction or later, two auctions or later from now, three auctions from now or later, four auctions from now or later?

These are (inaudible) going to consider them one at a time. The important controls in this regression besides fixed effects for types are going to be the number of unique bidders. I don't observe that, but I'm going to hope that putting in the number of unique bidders is going to (inaudible) nicely. I'm going to run it one order at a time, all the first bids, all the second bids.

My data sets come from the time when eBay was much more forthcoming with data, and so, I have the first bid and the second bid and I know everything about what everybody bid. These are 2001 and 2002. Types are titles for my DVDs and brand model combinations for MP3 players, something like Diamond Rio 500. I split the players into two groups because the prices vary a lot. So, I have low-price players and high-price players.

I've already said that I don't observe the number of unique bidders. Another big weakness with this kind of data is that you don't know what is actually being sold. Does this have some sort of a funky accessory? How used is it really? That's tricky. The only thing I'm going to put in is some control for new. Does the seller say it's new or mint or something like that? But other than that, I don't know too much about it except what was (inaudible).

So, I have all these regressions, two different orders (inaudible). Three data sets effectively after splitting the MP3 players. Three specifications of these forward-looking variables. There is preliminary evidence that most eventual winners only win one unit, so this unit demand makes sense. A substantial number of bidders actually participate in multiple auctions, however, which also makes sense. This model would predict that, in fact, you come in, you just lowball the bid, just keep lowballing it until you win at the right level. The level is determined by your valuation, your patience, your assessment of the competition.

So, even though in an isolated sealed bid second price auction you would have dominant strategy to put a value, here you have to assess your competition because tomorrow's surplus depends on how many people you're up against.

Finally, it does not seem that the model auction bidders simply submit a very low bid, then learn about something, and then later raise it to their true willingness to pay. If I look at bidders who bid on exactly the same thing twice, only 49 percent of the time do they actually bid a higher second bid in movies, 59 percent in MP3 players.

Another explanation I can rule out right away is the dangerous one. Since I'm looking at order statistics, it could be that these bidders just sort of show up and there's all sorts of cameras and they just close their eyes and pick one. Surely, if there are lots more cameras being offered in a particular hour, there would be fewer bidders per auction. If this sort of simple myopic process is happening, the order statistic of the bids would be mechanically related to future supply. In some sense, I would interpret all that supply as just future supply from the position of view of the early auctions. That is not happening because I control for the number of bidders as much as I can. I try both non-parametric controls, and I concluded that actually they are so nicely along a line that it's put in a linear effect.

You can see the results there. The news is good for efficiency of the market; good, in that all these categories, either type-specific or type-independent measures seem to be picking up some variance and they seem to be big effects.

What do I mean by big? Well, for example, in DVDs and movies, at an average price of $10, just having one of the next five auctions selling the same title reduces the price by 3 to 7 percent. In MP3 players, there is a similar percentage of price reductions by increases in the upcoming offerings. The type independent ones are insignificant in movies, and significant in the players, but not very big effects. So, in some sense, it seems that the type-specific effects are the big ones.

I also find that the second-highest bids tend to exhibit, percentage-wise, bigger effects than the first-highest bids. I'm thinking that it's probably because there's some remaining product heterogeneity which is not captured by my discrete types and also by the fact that my highest bidder cares so much more about that particular camera that he's not so susceptible to the future. I hope that makes sense.

To summarize, forward-seeing effects, like the ones that were captured by my model, do seem to operate on eBay. We're talking about three to seven price reductions when the same type is available within the next five auctions. Now, this is a very small amount of time, usually no more than a couple minutes. This is a fairly high lower bound on bidder sophistication. Bidders seem to be looking forward and taking that into account roughly consistently with the theory.

This provides direction for more fine grade structural models, and it provides cautionary notes for people who like to estimate levels of demand and interpret individual eBay auctions as independent. This shows that you can no longer say, "I have just observed 1,000 independent draws from auctions. Let's interpret that as 1,000 single-shot myopic auctions." No, unfortunately that's not it, and it would lead to a downward bias.

Now, some directions of how to think about this. Well, there was one hidden assumption throughout this whole talk that the seller is exogenous. The sellers are not exogenous on eBay. I'm not too worried about that, but whether or not there's going to be more cameras a week from now surely is going to depend on the kind of prices that the seller is going to observe today, and so, sellers may want to take note. The bidders are forward-looking. They are taking supply into account, and so, they may be best responding to future supply. So, the sellers may want to limit future supply.

Of course, this has relevance beyond eBay. Most sequences have look-ahead pre-announcements of this type. Procurements of construction have to have it by law. It's very difficult to run a sequence of auctions and keep the next item away from the bidder. So, these kind of models are likely to be relevant elsewhere.

Thank you.

DR. ADAMS: Thank you. To discuss the paper is -- Wedad is going to have to introduce herself because I don't know how to pronounce her name.

PRESENTATION DISCUSSANT -- PROFESSOR WEDAD ELMAGHRABY

PROF. ELMAGHRABY: Elmaghraby. It's phonetic except for the H, which always throws people. Thank you, Chris, for inviting me. This has actually been very enlightening for me. I'm one of the few people I guess here -- Ravi and myself who are on the operations side, not an economist, not an econometrician. So, in that sense, I've learned a lot.

I do research in auction theory, mainly procurement auctions and it's more theoretical. So, I'll put a bid in for that at the very end, a pitch for that.

And I'm also -- as one of the previous speakers said, I'm not an eBay user. Sorry. However, I am married to one who spends about a good three hours every night on eBay. So, he is my own at-home field experiment that I get to watch. So, when I was listening to all the speakers today, I was going, uh-huh, yeah, that checks, yeah, the sniping, the misspelling. That's one of his favorite strategies to try to look for things. The bidding rings, interaction with sellers, interaction with other bidders, you know, mark-ups due to reputations. All of this checklist, yeah. That's the kind of comments he gives me as he's bidding away.

So, I had his experiences in mind as I was reading Robert's paper, which I enjoyed reading. So, let me actually find my slides. Sorry.

So, this was -- it was very interesting to learn that nobody had actually done any research before on forward-looking or forward-seeing bidding on eBay because, especially on the operation side, it's something -- there's a group of us in operations that are doing work in auction theory, and that is something that we address theoretically, how does sequencing of auctions, the heterogeneity of items being auctioned off sequentially, how does that affect how people are going to bid?

So, given that this hasn't been done, this is a really interesting research question. How much detail of available information about near future auctions do actual eBay bidders use? And so, Robert did an excellent idea of summarizing the main components of his model, rational consumers -- well, he didn't test rationality -- independent, private valuations, single-unit demand for a specific type of good, an important thing, bidders don't see past prices and I'll come back to that, and that there's a fixed number of bidders in each auction, although who is participating will change over time.

In the auction sequence of single-unit auctions for different goods, second-price sealed bid, and he does the study for MP3 players and DVDs, and he's solving for infinite horizons that he's date fitting.

So, I put mine in just to see what he was talking about, I put in Sopranos, and I was looking for Sopranos: First Season, and all of this came up. So, he's saying if a bidder is presented with a page like this from eBay, they're really going to focus on the near future things, looking at the stuff in the red. And when you look at this quickly, you say, gosh, there's a lot of things happening, a lot of activity on Sopranos, and if I took the time, some bidders may take the time and say, well, actually, only a few of these are really first season that I want. Some of them, I just looked at the total number of items that are showing up in red.

So, what's the prevalence? How are people incorporating this information when they actually place their bids?

So, his results, they do seem to engage in at least one form of forward-seeing, the waiting times until the next auction of the same type increase bids, and the impact of another offering in the near future decreases the number of intervening auctions. So, that's what I took as my three main take-a-ways.

As I spoke with some of my colleagues who do work in eBay auctions, this is a great start. So, what's next? So, I took it, you know, when my husband, observing his behavior. I think that ran counter to the way he bids the most was this -- assuming that my husband is not going to look at past auctions because that's not -- the first thing he does is he goes and see what's been happening on similar stamps or, you know, whatever it is he's trying to buy. Not only in the past, but also what's happening currently, concurrent auctions, because some bid activity does happen, or at least what's being auctioned by whom.

So, I'm curious as to how these do not see past price assumption and concurrent auction assumption, they're not happening simultaneously, how strongly that would affect your results. As I said, there's some work being -- one of my colleagues has done some work on inference and she'll be on the panel later on, and there's some work going on in the OM group on auctions. You might be interested in a paper by Damian Beall (phonetic) and Larry Wine (phonetic), on pooling analysis of two simultaneous online auctions. It might give you some food for thought.

My own interest in auction design has always been, okay, what can I learn about auctions and then how can I help advise the auctioneer on how to better design an auction. So, what would you say to an auctioneer or to a seller now?

Given your data, how should he better use this to improve his profitability? Depending on the particular -- if he has multiple units or multiple DVDs or multiple units of the same DVD, how should he space them possibly? If it's MP3 players, should he space the high -- should he sell the high-priced items first and then the lower-priced items? You know, how can he use that?

Then this idea of auction design dimension, something that a couple of previous speakers have mentioned, the big money out there in auctions, although it seems to be on the decline, is the B2B auctions. There are a lot of advices needed. And there's so much talent in this room on the econometric analysis of auctions, I wish I could just redirect some of that towards procurement auctions, because that's where a lot of, I think, very, very interesting questions are happening. You look at free markets, which was a card by (inaudible), vertical net, combined net, and then large companies like HP, GE, GM, all of these are running their own internal auctions. They're starting to funnel more money into it.

I started to do some work with a colleague at HP Labs and they've been running auctions for the last couple of years and they've managed to save a couple of million dollars. So, now, there's a big push to use auctions a lot in their procurement process.

Well, take all of your skills and help them. The problem is, I mean, they want to run auctions, but they don't know how to design them. How much are people actually going to do this forward-seeing, forward-looking, the suppliers when they're bidding? How should that influence how I sequence what I sell, what I should lot together? There's a lot of dollars at stake, but there's also a lot of very, very interesting research questions. So, I enjoyed reading the paper. Thank you.

DR. ADAMS: Thanks, Wedad. Okay, I think we have Axel up next.

PRESENTATION: eBay: DESIGN AND BIDDING BEHAVIOR

BY PROFESSOR AXEL OCKENFELS

PROF. OCKENFELS: I'd like to talk about eBay. We have eBay in Germany, too, and it's quite successful. I think the percentage of households using eBay in Germany is higher than in the U.S. I'd like to talk about eBay's design and how the design affects bidding behavior. Here, for simplicity, I focus on private value auctions; that is, on auctions in which bidders know exactly their maximum willingness to pay, which are independent from each other.

I'll talk about bid amounts in the single-unit auction, about bid timing in the single-unit auction, and I will also talk about more recent work on eBay's multi-unit auction, called, by eBay, Dutch auction. That's a different auction from what economists call Dutch auctions, but I guess in the end, eBay will win.

(Laughter.)

PROF. OCKENFELS: Okay. As we all know by now, eBay is a second-price auction, so you should bid your value, that is, your maximum willingness to pay. The main argument has been put forward by Vickrey, and we have extended this argument to eBay in our earlier work. EBay explains the economics of second-price auctions on their webpage and they come to the same conclusion, they always recommend bidding the absolute maximum that one is willing to pay for an item.

So, do bidders follow this advice by game theory and by eBay? We tested this in a field experiment. We asked eBay bidders whether they want to participate in an experiment, and if yes, we told them to go to a specific eBay auction. The winner of the auction received, say, 100 dollars, and everybody else zero. This way, we induced values. We could then see, how much each bidder bid relative to the respectively induced value. Usually, outside experimental control, of course, you don't know values. The main result of the experiment is that a majority of bidders actually bid values. The average bid as a percentage of induced value was 89 percent. However, note that we only have this information for losers, because we don't know the bid of the winner of the auction, the highest maximum willingness to pay submitted to eBay.

Now, this is quite good support for the theory, and it is also supported by similar experiments conducted in the lab. We copied eBay's platform in the lab, induced values, and replicated what we saw in the field in an even more controlled environment. In the beginning of our lab experiment, bidders bid a little bit less than values, but, in fact, on average, bids converged to 100 percent of values. Again, this supports economic theory, and at the same time it stands in stark contrast to what we found in eBay's multi-unit auction.

But before I show you our study of eBay's multi-unit auction, let me talk a little bit about bid timing, even though we have addressed this issue already a couple of times today. EBay does not only recommend bidding the absolute maximum that one is willing to pay, but they also recommend to do it early in the auction, and there are basically two reasons for this recommendation. First, eBay is not an English auction but a second-price auction, so the highest bid wins and not the last bid, regardless of the timing of the winning bid.

And the second reason is that late bids run the risk of coming in too late. It turns out that this is a significant risk. For instance, Esnipe.com reported a couple of years ago that 7 percent of all snipes placed through Esnipe.com could not successfully be placed. We also did a survey and asked snipers how often their snipes failed to be accepted. Snipers told us that this happened on average in about 10 percent of all attempted snipes.

So, sniping clearly involves risks. Still, we have a couple of papers, and there are in total maybe about 20 papers on the table now, arguing that there are good reasons for sniping on eBay, both in equilibrium in private and common value auctions, as well as out of equilibrium when you are playing against naive incremental bidders, as David has told us earlier today.

Now, eBay Germany doesn’t seem to like sniping. For instance, it is not allowed to use artificial sniping agents like Esnipe. The basic argument is that it gives an unfair advantage to those who use these sniping agents over those who don’t use them.

On the other hand, eBay offers kinds of sniping tools themselves. For instance, you can let them call you at the last minutes of your auction, if you have been outbid, on your cell phone and then you can snipe via your cell phone.

(Laughter.)

PROF. OCKENFELS: Anyway, if you don't like sniping, there's a simple design solution to it. We have shown in a series of papers that, if you don't use a hard close but a soft close like Amazon auctions do, then the strategic reasons for sniping vanish. What is a soft close? In a soft close auction, the auction is automatically extended whenever a late bid is submitted, on Amazon by 10 minutes, on Yahoo by five minutes. I didn't really understand the argument that there are legal reasons not to use it, brought up earlier at this conference, partly because Yahoo is using it and Amazon is using it. But anyway, if you don't like sniping and if you believe sniping occurs because of strategic incentives, then theory suggests that you should use the soft close.

Now, this is theory. How does it look like in practice? Here's a natural experiment, where we observed the timing of bids on eBay and on Amazon. What you can see is that on eBay, there's a lot of sniping activity, and on Amazon, there's basically nothing going on at the end.

There are many papers now that deal with sniping. But they sometimes differ by how much sniping is actually observed. Just today, when I came here, I read a paper that says there is, say, three times more sniping in Sweden than in the U.K. There is all kinds of empirical work like this. But I think almost all these papers basically agree that there's a lot of sniping on eBay, but not on Amazon-kinds of auctions.

I should also say that more experienced bidders on eBay bid later, while on Amazon we see just the opposite effect. This supports the view that the strategic incentives for sniping drive the observations. It's not so much that the naive irrational bidders are those who snipe. The sophisticated bidders are those who snipe on eBay.

Now, we investigated this also in the lab because, there are many things that are different on eBay than on Amazon that may partly account for our results. There are much more bidders on eBay than on Amazon, et cetera. But in our lab experiments, the only difference between eBay and Amazon is the close of the auction, hard or soft close. However, in the lab too, there's much more sniping on eBay than on Amazon, and the effect is increasing with experience. Since this is exactly what we see in the field, the lab results imply that the different rules by which the auctions end alone can organize the patterns we see in the field.

Now, I want to make clear that this research does not imply that we should recommend eBay to change their ending rule, because there are other issues, like the entertainment value of the hard close. In Germany, they have these advertisements on TV and everywhere, which focus on the excitement related to all the bidding activity close to the hard close. So there are other issues. But again, if you don't want sniping for any reason, there is a simple design solution. This is one of the points of the research.

Let me now finally come to bidding in eBay's Dutch auction. There's hardly any study on eBay's multi-unit auction, which is somewhat surprising to me because, obviously, multi-unit auctions are important. Of course, we can think of eBay as a huge multi-object auction with millions of objects sold simultaneously and sequentially. But here I am focusing on the design of eBay's multi-unit auction format.

If you want to sell many items at the same time on eBay, you can use what eBay calls the Dutch auction format. As a buyer, you have to submit a maximum bid along with the number of units that you want to have. EBay then explains that all winning bidders will pay the same price, which is the lowest successful bid. Now, is this a clever design? We all know that designing multi-unit auctions is much more complicated than --

(End of Tape 4, Side A)

PROF. OCKENFELS: -- single unit auctions. In fact, even in the very simplest case of an eBay Dutch auction, namely if all bidders demand at most one unit, there are problems with the design. Why? Now, a natural extension of eBay's single-unit auction would be to have a uniform price for all winners in the Dutch auction equal to the highest losing bid. Then, you would have the same incentives as in the single-unit auction: you should just submit your maximum willingness to pay. But in eBay's Dutch auction, the final price equals the smallest winning bid. So, because a winning bid determines the price, there are incentives for bid shading. You should not just submit your value.

~ Private information redacted at Speaker's request ~.

So, the conclusions are, first, with respect to behavior: incentives matter. Bidders respond to economic incentives. Now, they don't do so always in a rational way. We don't always see equilibrium behavior. For instance, I think that one reason we see so much sniping on eBay is that sniping is a good strategy in equilibrium, but also out of equilibrium, against certain kinds of common naive bidding behaviors.

While bidders are not always rational, they respond to incentives in a systematic and a predictable way. This opens the door for scientific market engineering based on field experiments, lab experiments, and theory, and this is why we can help eBay and other market architects to optimize their auction and market platforms.

With respect to institutions, the conclusion is that details matter. Even small details in the design may have a significant impact. For instance, what happens at the very end of the auction, in seven days, may have a huge impact on whether and how people bid early, it may also have an impact on efficiency and revenue. Or seemingly little things like whether you have a price equal to the smallest winning bid or equal to the highest losing bid may have a big impact, as I've shown you in our field experiments.

Thanks a lot.

DR. ADAMS: And David Porter from down the street, GMU, is going to give the discussion

PRESENTATION DISCUSSANT -- PROFESSOR DAVID PORTER

PROF. PORTER: Thank you so much. My discussion is going to be fairly short because David Reiley did a good job of discussing this paper previously. But let me tell you about my center of gravity. I do controlled laboratory experiments. So, when I was reading this paper linearly, going from the front to the back, the first thing I see is, well, it seems weird. Standard theory says they should bid their value and we don't see that, and I thought, hmm, welcome to the club. I see it all the time.

Second is, I see there's a note that says, eBay recommends that you bid your value early, and the cynical side of me said, that's the reason why people don't do it.

(Laughter.)

PROF. PORTER: Okay? So, anyway, I'm reading the paper and I get to a very interesting part of the paper that wasn't presented here that had to deal mostly with common values and dealers in antiques. And the part in there is that they care which auction they're in. They would rather be in an eBay auction than the soft close auction so that they don't have to give away their information, which is value.

And then I remember a paper by Dan Houser, it was a field experiment in which he used Yahoo and said you could choose the close you wanted, soft or hard, and so, he ran some experiments where they used gift certificates, so it was pretty clear sort of what it was, and he found out that if the sellers used a soft close instead of a hard one, you make more revenue.

So, it dawned on me that, hmm, there's something missing in what Axel was doing and that is, it's not just the buyers, it's the seller, too. They make a choice. And then there's, I noticed, another paper which is based on a very simple, oh, comment made by Vernon Smith to me. He said, you know, in an auction, the most important person to show up is the second-highest value guy. That's where you get the price information, that's where you get the revenue information. So, participation is kind of important. And so, there's a paper by Roman Bragoff (phonetic) who ran some experiments in which the seller can choose the type of auctions he wants, and the buyers show up and they pick the auction they want to be in.

After about 20 times, everybody's starting to pick eBay auctions. Why? Because even if you have a second-highest value, you've got a chance to win. So, you show up. If every time you show up, you lose, why show up? It's costly. It doesn't make much sense. So, in a sense, I don't think we need to tell eBay how to design and run their auctions, I think they know what they're doing. They get participation, they get people showing up, they get the buyers to show up. When the buyers show up, where are the sellers going to go? You can go to a place that gives you higher revenue, but if there's no buyers there, why show up.

So, in a sense, you need to put the sellers in to see what the equilibrium is in these sort of things, right? It's not just the buyers, although I can see why you'd want to focus on that. I guess that's all I have to say.

DR. ADAMS: Thank you, David. Hal, I'll let you ask a question.

PROF. VARIAN: Okay, here's the question. So, as I understand it, in Japan, the Yahoo auctions are larger than eBay auctions, I believe.

UNIDENTIFIED MALE: I don't know.

UNIDENTIFIED MALE: (Inaudible)..

PROF. VARIAN: Well, it's true.

(Laughter.)

PROF. VARIAN: So, the question is -- but at one point you did, no?

UNIDENTIFIED MALE: We did and (inaudible).

PROF. VARIAN: Yes, okay.

UNIDENTIFIED MALE: And Yahoo was much more successful.

PROF. VARIAN: Right. But does Yahoo use the hard close or the soft close in Japan?

UNIDENTIFIED MALE: I don't know, but at one time, you could choose the one you wanted.

PROF. VARIAN: So, the question is, is it really a path dependence issue, because after all, the buyers aren't going to show up if the sellers aren't there, the sellers aren't going to show up if the buyers aren't there. So, is it really just -- you could do it either way, it depends on which one starts first and gets the network effect going. I mean, that's a conjecture. I'm just --

UNIDENTIFIED MALE: I don't know.

PROF. VARIAN: We don't know. We'll try to find out.

UNIDENTIFIED MALE: It seems like I need to run another experiment.

(Laughter.)

DR. ADAMS: And next, and the last sort of paper presentation, is Robin Sickles from Rice University.

PRESENTATION: ESTIMATING CONSUMER SURPLUS IN eBAY COMPUTER MONITOR AUCTIONS BY PROFESSOR ROBIN SICKLES

PROF. SICKLES: We lost, Chicago won.

(Laughter.)

PROF. SICKLES: What do you do?

This is joint work with Tugba Giray , whom I've never met; Kevin Hasker, whom I've seen most recently about three years ago when he was -- maybe four years ago when he was still at Rice; and myself. We started this project in 1997 and spent several years putting together software that would allow us to get the data. I'm going to talk a little bit about that in a minute. And then various people in the original research group that involved Kevin and a former Ph.D. student of ours migrated to Mexico and to Turkey. So, I'm very glad about this conference because it did sort of provide me with some leverage to get back to this research agenda.

I must say that a hard-stopping rule of today at 4:00 or 5:00 did get me to this stage, which wasn't at all strategic.

Everybody knows that eBay is a big player. We don't need to spend a whole lot of time motivating why studying eBay is of import. I will mention, though, that when we started this project back in 1997, my interest was in looking at consumer surplus and trying to figure out to what extent these mechanisms generated consumer welfare. And sort of an interesting isomorphic problem to looking at auctions and looking at order statistics and all of the measurement issues that go with trying to identify moments of distributions from order statistics is a literature in stochastic frontiers where you're trying to benchmark firms and I've done a lot of work in that area. So, there seems to be a tremendous amount of overlap in terms of the -- certainly, the generic estimation methods and that's sort of where I come into this forum.

So, consumer surplus is what this paper is about. What we're trying to do is to estimate a structural model where that structural model will lead us to a vehicle for identifying consumer surplus in eBay auctions, these are eBay auctions for computer monitors. We're going to be estimating bidders' values. We have an entry process as well. We're going to be using parametric methods. Again, one thing that I've discovered from a life of working in estimation is how squirrely results tend to be when you're basing them on order statistic. That goes sort of even more -- it's even more the case that the results are difficult to get when you're going to the non-parametrics. The approach we're taking here is to look at a variety of distributions for the private values and to utilize standard specification tests, likelihood measures, to extract, at least, an optimal set of results from particular distributional forms.

So, we're going to test for sensitivity of the distribution, we're going to test which distribution fits the data best, we're going to test how well each estimate performs against a non-specified, non-parametric distribution.

We have about 3,000 PC color monitor auctions. The screen sizes go between 14 and 21 inches. They were auctioned between February 23rd, 2000 and June 11, 2000. There have been -- and that's what we're going to use to get these estimates of consumer surplus. Now, you know, there are others that have looked at consumer surplus, one of whom is going to be talking shortly, Ravi. Quan Vong has as well. In fact, in the original work that Gonzales and Kevin Hasker and I did, we estimated consumer surplus of eBay auctions based on the Quan -- the Jean Jacques Lafont/Quan Vong and their student, their paper on eggplants (inaudible), the simulated non-linear least squares methods. But we didn't pursue that as much. But there have certainly been others who have done that recently.

You all know about the common format for the eBay auction. We're going to be looking at single-price auctions, which compose about 87 percent of our data. We're not going to be looking at the multi-unit platform auctions.

We base our techniques on the methods that were developed by Donald and Parsch. We don't have to have -- there's some funky problems with the regularity conditions for maximum likelihood that sort of justify moving to other types of estimators, but those, I don't believe, are at play in the kind of data we've got, at least here at this point.

We do have in the eBay auctions auctions where no one decided to bid. So, we do have a full auction, it's not just auctions where the unit was sold.

There are, obviously, other approaches, semi-parametric approaches. I'll mention that I'm a bit skeptical just on the information content on estimates from non-parametrics of order statistics. But that's not to say they don't work or they can't work. It's just that I think you need a lot of information. The Song approach utilizes, I think, the Gallant (inaudible) Polynomial expansion of density function and uses the difference between second and third order statistics to tease out actually the dependence of the likelihood function on number of bidders, potential participants, I should say, not the number of people who have bid. There's the Bayesian methodology obviously by Bajari and Hortascu, and then the Lafont -- Ossard, (phonetic) was the student and Quan Vong's work.

Just very briefly, we do have data. We have a particular protocol I guess the lawyers have left, so I don't have to worry too much about this. I was told that this was not actually the best thing to be doing, but we did it anyway. Well, we did it before I realized that there was an issue. I guess there are several lawsuits out there which I was made aware of actually by Kevin recently. But I do need to talk to the eBay people to see if this is still something we could do because we spent a lot of time putting together this software. We have a spider program that periodically searches eBay for recently closed, in this case, computer monitor auctions, downloads pages, gives item and bid histories.

We've developed a software program done in Python, which is this multi-platform, multi-OS project-oriented programming language, and we have three parts to that data of access protocol. that are -- you know, we go out and grab the data, we parse it at the web pages and the HTML code, and then we iterate through database entries to create a tab delimited file, and that gives us information originally on about 9,000 auctions.

We've needed to get that down to a smaller number because we had to match it up with additional data and we had a bunch of further data processing with some of the other raw HTML files. We did string searches to collect extensive descriptions of the entire data sets. We had characteristics of the data that I think are fairly straightforward, but also difficult to sit there and code,you know, in a hand way.

This is a way that you could develop very large data sets, data sets as large as you want, of the variety that Hal was talking about, but with variables, with a substantial amount of co-variant information about the auction -- about the auctioned item. We ended up with about, in this case or this set of estimates, we're looking at 29 to 34 auctions. And I mentioned that we're looking at PC monitors. We have a variety of bid retraction cancellation criteria that we have to address, but I'm not going to talk about that, we don't have the time. I will mention very briefly what the -- how about if I do -- can you see that okay?

This is the data or at least these are the descriptive statistics for the data. We've got size of the monitor, we've got the dot pitch. You know, if the dot pitch isn't available, we dummy that. We've got the resolution of the monitor or dummy if we don't have that information. We know whether the monitor's new, whether it's at least advertised as like new, it's been refurbished, whether it has warranty. We have brand name, whether it's a flat screen. We also have the seller's feedback rating, length of auction. We have dummies on the sizes of the screens, whether there's a secret reserve price, whether it's been met. We also, obviously, know what the price is, if it's not -- or what the reserve price is when it's not secret. Number of bidders and I just focus on one -- well, let me focus on two things.

One is the mean of the number of bidders, which is around four, and the other is the -- actually the median. Let's look at the median. The median of the number of bidder is around three and the median of the sales price is about 100. The reason I'm going to look at medians is because these estimates are from likelihood functions that are very, very ill-conditioned to say the least, especially the structural models.

So, you've got ratios of all sorts of crazy stuff, and the moments don't exist in finite samples for those probabilities. The summary statistics themselves are themselves just ratios of a bunch of estimates, and again, those don't have moments. So, you know, we could have just truncated the tails. We didn't do that, we just calculated medians. Obviously, the more appropriate way to do it would have been to trim, but we just are going to be looking at medians. So, all this stuff I'm going to be talking about in the next few minutes are just medians, okay? We do have means.

If you look at the paper, which I'd rather you not do until I can get a newer version of it, but you'll see values of estimates and they're bouncing all over the

place. The medians, however, actually are rather robust and rather informative, I think. Of course, I thought Houston was going to win. I was at the 18-inning game, too.

Okay, we've got more on the string searches used to construct the data that I just looked at. We've got some -- what's the model? The maximum likelihood or the likelihood functions. We've got a standard sort of specification of the winning bid. We've got a private value, a second-price auction. We've got a number of participants which, itself -- we're not talking about the number of bidders, but the number of potential bidders because some of the people don't bid. We have CDFs for the bidders' values, PDFs for the bidders' values, dummies that indicate if there's no participation or if there is one participant, this is a second-price auction. That's the likelihood given the number of bidders. I'm sorry, given the number of participants.

We're going to allow for entry and entry to have a standard price on entry process with intensity parameter lambda. If we enter or if we bring that into the analysis, and obviously, the length of the auction is set, so we know what the length of -- you know, what T is in this (inaudible) process. The likelihood for a particular auction then is going to be the conditional times the margin and we have our total probability. And that's going to be maximized over all of the 2,953 observations.

And there's some intuition, you know, behind what that stuff is. But, you know, I don't have the time to spend on that.

I will talk a little bit about some of the results, though, and what I might mention as well are results across these different distributions. Again, I'm not going to be looking at a non-parametric estimator that would, you know, presumably nest these, but different distributions, all of which have this similar kind of pattern or at least could have a similar pattern. Certainly, the gamma is more flexible than others and we'll find that, duh, it is the one that fits the data the best. But these all have , you know, similar kinds of shapes. And, in fact, when we look at the testing results, we'll see that there's not an awful lot that discriminates these other than the fact that the gamma seems to be the best and, I might mention, that the results for the gamma are really the most reasonable.

There's a lot of stuff that's going on in this process of estimating these parameters. There's an enormous amount of data collection and automated data protocol. There's a substantial amount of programming effort. There's a substantial amount of, obviously, mistakes made before we got this thing right. But at the end of the day, you'll see, at least I hope you'll see, that the results are turning out to be rather reasonable.

Let me go back to this other file real quick. I'm from Pittsburgh. I never said file before I went to Texas.

These are the results. Just to let you know that I actually did calculate them. Nobody's going to look at those, but those are the results for all these parameterizations of the private values. They're actually rather stable, okay? And there are some issues we can certainly pursue, maybe after this is over, on the reputation issue, which I think is one that does have some interesting implications from our work and that of others.

Let me mention one thing. Remember the value or the median price for the monitors was 100. The median price is estimated to be 93.1 for the gamma distribution. Actually, let me do something first. If you look at the different criteria for the evaluation of the models from the different distributions, using AIC, BIC and BCC, as well as a value of the likelihood function, since these don't have terribly different numbers of parameters -- I mean, there's only a couple of parameter differences between these -- you know, with a shape and a scale parameter, you wouldn't expect the results to vary that differently among AIC, BIC and BCC relative to what the value of the likelihood function was, and they're completely consistent. The gamma is the one that, at least in terms of a blind man statistical criteria, is going to be viewed as the best.

So, you know, let's just go back with that fairly straightforward automated result and look at some of the summary statistics. We've got an intensity parameter, lambda, that's estimated to be about 10. Now, that's going to be the average number of arrivals, the average number of --

UNIDENTIFIED MALE: (Inaudible)..

PROF. SICKLES: Yeah, yeah, that's what I'm trying to say. Thank you, thank you. It's always nice to have my prose edited, but you're a good guy and I know that.

The average or the expected number of participants in the auction. You know, we had about -- six was the median. We had a potential number -- that was the median observed. We have 10 as the median for those that came into the auction, did bid or didn't bid. Those are actually a lot more stable than I would have thought across the distributions.

We have a consumer surplus of -- then we had this predicted value of the computer monitor for the gamma that's maybe about $7 less than what we observed. We've got a value of -- let me just back up a little bit and look at how we construct consumer surplus. We've got -- in terms of just the normal market, ex post consumer surplus, which is, obviously, the area underneath the demand curve and above the sales price. We can actually get that ex ante and ex post. Ex post, we can get it without having to worry about the number of potential participants in the auction. And there's some not terribly heavy lifting that gets us to that conclusion.

This is actually then a result that one could compare to Song's results, because Song's methodology, where she uses the difference between the second and the third order statistic. , and then there's, you know (inaudible) and his student (inaudible) approximation. It's basically like a non-parametric MLE except it's using a polynomial series -- a Hermatian (phonetic) polynomial series-- to approximate the density function. In that particular case, it's almost like parse of likelihood in a duration model. You don't have to worry in the parse of likelihood about the baseline. In her estimate, you don't have to worry about the number of participants. And in this calculation -- and hence, her calculation for surplus, based on her non-parametric method, doesn't require specification of the number of participants. Here, the consumer surplus, if it's ex post, doesn't either.

We calculate that value. Again, it's independent of I. We have little proof of that. We also can construct a lower bound for that consumer surplus. There's a consumer surplus measure where we essentially assume that the first and the second valuations are the same in every auction. I was constant, it was endogenous, then this would be a precise lower bound. What it does is give us an estimate of the consumer surplus that is independent of the tails of the distribution. The tails of the distribution, when you're looking at a variable that's got support on the interval that's bounded by the number of observed bidders and positive infinity -- obviously, it's truncated -- you know, can give some squirrely numbers depending on what those tail properties are. This, in the sense, is independent of those tail properties.

So, if we then go back to these estimates of consumer surplus, we're looking, again, at the value of the consumer surplus expected. Now, this is, again, ex post., from eBay and then the lower bound in that table six of estimates of consumer surplus. Now, recall that the price was 100 and go back to the summary statistics. That gives us a lower bound of about 26 percent and an upper bound of about 49 percent, 48 percent consumer surplus in eBay auctions, which is not a small benefit to the consumer, this kind of mechanism or mechanistic design. Actually, it's 42 percent, I should say, I'm sorry.

There are methods that we developed to find the best distribution. There are both parametric, as well as non-parametric ones, the ones that are based on likelihood, I've already shown you, and there are tests, as well, against the non-parametric distribution of the third-highest values, which we don't explore at least yet, but we certainly will as this paper gets more fine-tuned.

So, that's that. I don't have a slide that says the end.

DR. ADAMS: Thank you, Robin. We're going to have Ravi Bapna from UCONN to talk about this paper.

PRESENTATION DISCUSSANT -- PROFESSOR RAVI BAPNA

PROF. BAPNA: Thank you. It's fun being the last discussant for the last people. Everything's already been said. I really have nothing new. It also actually makes things a little easier because a lot of things have been explained already. So, let me, you know, begin by talking about the -- some really, really strong points about the paper. I mean, one thing is that it is actually a paper that they're currently working on. So, you know, there are a lot of versions of the paper over the last week or so, in fact. So, it's been interesting to see that.

Overall, I think, the topic is very important basically because it maps consumer's welfare to dollars. So, you actually get a sense of how much in dollar terms people benefit out there, consumers at least. And if you think of it from a policy perspective, that's really important. You know, if eBay switches to, let's say, a different bid increment policy or it does something else, and if we have benchmarks, you know, through Robin's work and other people who are doing work in this area, then we can go back and actually measure in dollar terms what the impact of that is to the consumers. So, I think that is very important.

It's also -- it's something that's not, you know, very widely reported. I mean, there have been some other studies in non-electronic markets that I've looked at it, but it's a tough thing to measure, especially if you think of posted price markets. It's a tough thing to measure. When you're all ready to check out at Best Buy and you buy a DVD player, they don't really ask you how much were you really willing to pay for this, you know. So, that's why I think it's a tricky thing to measure and it's a very important topic if you don't have a consumer surplus index or any such thing. So, I think that's something that -- you know, I really enjoyed reading the paper actually.

The other really nice feature of the model itself was that it actually models the zero bid case. The fact that there are a lot of options on eBay, and people don't talk about it much -- a lot of options on eBay don't get any bids and this is some cost to the seller. You know, what I think is seller's lament. And we try to actually ask people at eBay whom we knew and for whatever reasons, obviously, and luckily, I guess, they don't really disclose that. EBay gets their commission, but the sellers essentially don't get a sale out there. So, understanding what -- you know, I mean, obviously, you know (inaudible) the second-highest valuation is important, but I think the first-highest valuation comes before the second highest. So, you need to know what -- how do we get the first person to come and bid on that auction, and understanding that, I think is very, very nice. So, that's another thing that -- since it's a part of the model and the estimation, it was new to me, at least. I hadn't seen that in other work.

Let me talk about the empirical part. I think the data is very rich. It's very well passed out, very clean looking at the actual, you know, characteristics of the item and, you know, that's a part of the model of sort of the likelihood of observing that very low auction. So, it gives you weights. Ever think of, you know, how much weight is a 17-inch monitor in terms of a preference versus a 19-inch monitor and things like that. So, that has a lot of implications to, you know, retailers who are designing these products. You know, what are the actual sort of values consumers are putting by working with their dollars? This is not a survey. We're not really asking them, do you like 19-inch versus 17-inch or whatever. These are weights that we're estimating in very good situations, basically. So, I think they're very reliable.

The bid is homogenous, and I'll get into this a little bit, in the sense that it's looking at computer monitors, which is nice in this case, at least, has a lot of value, and very extensive. It's a very large data set. It's a very impressive job of actually going out and collecting it, and, you know, having been involved in this kind of stuff for a long period of time, I totally understand the effort and the sort of trials and tribulations of doing it with the eBay (inaudible) and (inaudible) interface and all kinds of stuff. I mean, a lot of people (inaudible) relate to that.

The other really nice thing about this paper is that they don't assume, you know, some given distribution. They actually -- so, it actually methodologically is nice because if you really want to -- you know, it specifies how you would go about doing this in five different distribution (inaudible) and then to try to find the best fit. So, that was another very sort of interesting and nice feature.

Then looking at the results, I just have the lower bound right here, you know, the fact that consumers capture at least 26 percent of the total surplus, that's a new and good finding. I mean, it's really important, I think.

So, with that -- I think, you know, as I said, it's a work in progress, and Robin's been very kind in sort of keeping me up-to-date with everything and answering some questions over the last couple of weeks. So, I'm looking forward to seeing more. Let me also actually -- and in the paper, actually, they have been very open about asking for suggestions. So, let me give some humble suggestions here.

With respect to the entry, Galit, who's going to come on the panel next, actually has what I view as a richer specification of the arrival in eBay. It's actually a non-homogenous process. Basically the idea here being that the different stages of the eBay auction -- everyone talked about sniping and so on -- they actually are different, you know. So, the arrival rates are actually different in different stages. Galit will talk more about that. I don't want to take her thunder here. She fashionably calls it a barista process. So, I'm going to let her spell that out.

That's -- Robin didn't have the time to get into this, but looking at the results, actually, the reputation of the seller rating (inaudible) for experience and reputation. So, on the regression when they're actually predicting -- when they're looking at the value of the item, they get a negative coefficient with a seller rating and I believe it's not significant. But on the entry -- on the entry estimation, they actually get a positive coefficient. So, this is, I think, again so those of you that are looking at the reputation literature, it seems to suggest that, you know, there is something more than just simple, you know, sort of -- you know, better sellers get better things out here. I mean, I think what's really happening out here is the people coming -- once they come into the auction, then the effect of the reputation sort of disappears in some senses.

But still, I think -- I'm still a little bit troubled just looking at, you know, the negative sign out there. So, my suggestion would be to, again, people already talked about this. But my suggestion would be to break up the ratings that you have by the ratings that were obtained as a buyer versus ratings that were obtained as a seller, and also break up the positive neutral and negative neutral I think weighs as much as a negative in my -- it's a big statement anyway. So, I think you've got sort of two or three kind of more options out here to look at the rating and see whether, you know, it plays differently maybe.

So, I think there's something more that you can probably play around with. There are some variables that I didn't see for whatever reason. Maybe I missed out on something, but the opening bid itself, I think, has been shown by Ali and Patrick to influence entry and so, this was something that maybe I missed out on, but it should -- I think it should be there, along with the other sort of variables that you have there. And perhaps, also, with throwing in the shipping costs, okay? There is a school of thought that believes that sellers really, you know, make money on the shipping. So, to the extent you can control for that. If you can just throw it in (inaudible) and what was the value that would be nice.

One last quick comment, okay, I'm going to skip this one. They have really interesting things to say about our work, which is very similar, you know, doing the same kind of thing, but in a totally different approach. So, the words that are used out here are brilliant, unfortunate and (inaudible), okay? So, how do you put this together?

So, my sort of take on this is that if you look at the very front page on eBay, it's called Listings on eBay.com, okay? This tells you basically in a snapshot all the auctions that are going on on eBay right now. If you add this up, I did this yesterday, it's in the order of three million, okay?

So, so, the work that Galit and Wolfgang(phonetic) and I were also (inaudible) said something about the market as a whole, okay, and I think Ali's paper in the morning was the other (inaudible) that looks at all the different categories. And let me just take half a minute or two and just sort of quickly give you a sense of what we're doing out here.

Basically, David talked about E-Snipe. I have a bidding agent called Cniper, it was actually free, so people go to Cniper and a lot of people from Germany, believe it or not, go to Cniper.com, Cniper with a C, and they go and they bid on eBay using this agent. So, for those people who win the auctions using this agent, I actually know the highest bid, okay, and then, you know, I know the price and I know -- maybe know the other stuff. So, basically, it's a totally different way of estimating surplus, and to the extent, you know, you can say (inaudible) value; you can say it's, you know, sniping and you can pay a second price (inaudible) at auction, theoretically, you know, you're in good shape, you can basically rely on (inaudible).

If you get into common values, okay, then, then there's a whole estimation that you -- which I don't have the time to get into now, but, you know, we have a copy of this paper, so, you know, happy to share that with you, but overall, I think I really, really actually enjoyed the paper and I'm really looking forward to seeing the final version of it.

DR. ADAMS: Thanks, Ravi. In the world of very cool data sets, Ravi has the coolest.

(Laughter).

DR. ADAMS: Let's bring the people for the last panel. Before I do that, could some of the RAs stand up who have helped us. I want to just give them a round of applause.

(End of Tape 4, Side B)

(Applause.)

DR. ADAMS: Okay, can we -- while we set up, I'll give Robin the mic.

PROF. SICKLES: Just in response to one thing that Ravi said, this was a joint paper. The term “unfortunate” was added by Kevin; “brilliant” was added by me. I want to make sure you understand that.

(Laughter.)

PANEL: INFERENCE FROM INTERNET AUCTION DATA

DR. ADAMS: Thanks, Robin. First up, we’re going to have Galit, Ravi’s co-author.

PANEL PRESENTATION BY PROFESSOR GALIT SHMUELI

PROF. SHMUELI: Hello, I’m extremely happy to be here today. I’m somewhat of an unusual person in this crowd, and although I’ve been working in the field of auctions about three-and-a-half years, I am actually a statistician. I am not an economist and I actually started working in this field inspired by some of the early papers that I’ve seen on David in 1999 and I always wanted to sort of get into auctions since then, but, you know, I was at Carnegie Mellon for a while and I said, hey, do you want to do auctions? And they said, nah, not interested. And then I moved to the University of Maryland and there was a very strong group of people doing auctions there and that’s what kind of -- I said, I‘m going there.

Then I took out my little sticky note with auctions and I said, does anybody want to do auctions from a statistics point of view, and luckily, I found Wolfgang Yunk (phonetic), who’s also a statistician there and my colleague, and we’ve decided to actually tackle a lot of the issues that were going on here in this wonderful empirical world through a statistics approach.

So, there's been a learning process absolutely and I've learned a ton from a lot of people, a ton from Ravi, and what I'd like to say today is, again, it's going to be very different, I think, from a lot of what's been going on here because it's from a different set of eyes. But to show you a little bit about how I'm thinking about things and what are some solutions to some problems for a lot of these issues that were mentioned today, and this is just a completely different approach.

So, I'll have three main themes going on. One is about sampling and, again, you know, I'm not afraid of the web crawling, fine, put me in jail. But that's what everyone's doing, so we should at least acknowledge what we're doing and how it relates to the estimation. Then I'll make a point about visualizing the data and how that's been extremely neglected and how insightful it can be, and I'll just show a few graphs of what's going on there.

And, finally, this word "dynamic" showed up in a bunch of different talks, but how do you think of dynamics, how do you measure dynamics, how do we integrate those into models? We have a really amazing way in statistics that's pretty new, also, to this field, but I think does an amazing job in auctions or in ecommerce type data in general.

I have to say that I was not planning to say a word about the barista, but now that Ravi actually forces me to do that, that was a work -- the barista is bidder arrival in stages, and we call it barista because we actually saw the interesting paper by Roth and Ockenfels on the self-similar property of the bidding and how the -- if you look at the bids in the last -- the whole auction, the last day, the last three hours, whatever, you see the same cumulative distribution until the last minute or two where it breaks down and it's kind of uniform. We said, huh, that's really interesting, but a lot of papers are assuming this Poisson process of arrivals and it's clearly not Poisson if it's self-similar. Let's try to tackle it heads-on.

And we created this really neat model with Ralph Russo (phonetic) from Iowa, a statistician as well, which is, again, a non-homogeneous Poisson model. We have all the code on the web and estimation is pretty straightforward. It works very nicely. It approximates very nicely the bids and how they come in in different -- these different stages. So, that said, I'm moving forward.

So, the first point is about sampling. Now, I think that unlike a lot of other websites that are kind of opening up to giving users data in a usable form, using X amount or whatever, a lot of the eBay analysis is being done by scraping the web, by picking up data ourselves, or if you're lucky enough, you have the field experiments going on. I know there are companies that buy data, but they're mostly in aggregate form and they're expensive for us, and et cetera. So, I'm just looking at what's happening now out there in the research community and most of us are actually picking up the data from the web.

UNIDENTIFIED MALE: Just one quick question on that. The API (inaudible).

PROF. SHMUELI: Yeah, but it lets you do for free 50 hits a day and we're talking about this large amount of data that we're taking off. So, I mean, if the

API would open up as big as like Amazon does, I think everybody would be very happy and we'd all be in better shape from a sampling point of view.

So, let me just make a few points about the sampling. So, in a very, very general way -- and, again, I am not one of these, you know, web spider whizzes. I just understand a little bit about how it works and can do a very primitive version myself. So, first, you get a list of the pages that you want to take data off and then you go and send the crawler to pick up all the data from those pages and put it into a database. Generally, those are the two steps.

Now, if you're doing that and you're going to make a lot of inference from that data set that you end up after some cleaning and whatever, there are going to be different issues. Some are just non-sampling errors that are going to be in there, things like -- there are two different types of non-sampling errors, things that are called measurement bias where basically, you know, if you're using a search in eBay and saying, I want to look for all the auctions, there's something behind the algorithm that's ordering them according to a certain order.

If you're thinking about data that get unrecorded for various reasons, you'll never pick it up no matter how you sample, right, because it's just not recorded. If there's a data refreshing policy, you know, if you're trying to get a date and it just refreshes it every other day -- that's kind of a crazy example -- then you're not getting the most refreshed data. There's a different type of non-sampling errors, which is called selection bias. So, if people are giving you untruthful bids or there's a poor website design so people are putting in erroneous things in there or -- and, again, this is more general than eBay, but it also applies to eBay. If your web spider interferes with the traffic somehow again -- and that might not be such an issue on eBay unless you're really bad, there can be robot restrictions right in the robot dot text that tell you don't go to this area, but go to this area. So, a lot of these issues will, you know, just give you non-sampling errors that, fine, you just have to acknowledge those and realize that whatever you're making inference at the end is going to be valid not to those areas and take them into account.

But then there are also issues that have to do with the sampling itself. So, right, I mean, we do have a fixed cost for writing this spider, and then you say, okay, so I can just pick up 50,000 more, it doesn't cost me anymore, I already wrote the program. But that's not completely true because in eBay, for instance, there are some categories that are harder to collect, where you need more log-ins, where you need whatever. So, you'll say, oh, let's just not bother with that category, let's just not go to adult material or eBay Motors because it's a whole different thing, let's just leave those on the side. You know, a lot of times, it's just a technical constraint, but then you have to, again, think about what that means.

Regarding a main theme, I think that most of the papers that I've seen, I think every one that I've seen, assumes that they picked up a random sample of eBay. So, they say, we've picked off all of the auctions in the last 30 days for this particular item. Some people assume that it's the population, but most people assume that they do have a sample because really the population of interest, even if you think you have all the auctions in the last 30 days for this item, is you're not making inference for these 30 days, right? People are talking about a year of surplus or a month or whatever period. So, of course, there is a population and, of course, there is reason to do statistical analysis and not just report numbers.

But furthermore, assuming that you have a simple random sample means that, in a classic way, you would have to do a list, like a phone book of all the possible auctions and then you would just take a random sample from there. That would give you a pure random sample. On eBay, when you're writing these web spiders, you don't have a phone book to tell your spider pick randomly these 20,000 listings and call them. That's not how it works. So, the problem is, if you don't have that, it's very hard to assure that your sample is really a simple random sample.

On the other hand, there are other more advanced techniques that people use outside of the online world, like surveys and anything like that, which are stratified sampling and cluster sampling, which really are helping you to avoid the need for this list. They kind of create a list on the fly. These actually turn out to be very, very natural mechanisms for ecommerce or, again, eBay, because a lot of the sites, including eBay, is hierarchical by nature. You have categories within categories. You have sub-categories. Sometimes you're interested just in one little item, but as we've seen now, I mean, we had 30 categories and Ali also had 30 categories and people are now comparing soft close to hard close.

So, all of a sudden, you are interested in a little more than just one little item, and then you have to start thinking, well, how do I sample from those two? Do I just take a random sample and see, oh, I had 20 percent soft close and 80 percent hard close or do you strategically go and collect data from those strata so that the inference will be valid? Or if you've already collected things, you have to think how to adjust all the biases that end up at the end of the day, and I think that can adjust these numbers that we're getting that can be, you know, huge differences if we correct for those biases. So, that's on the point of sampling.

This is just a tiny little example. I don't even think there's -- I'm just going to drop this.

Okay, on the topic of visualization, I threw down just a few graphs to show the value of adding graphics into the research that we're doing here whether it's fraud, whether it's mechanism design or whatever it is. And an interesting thing is, when I talk to statistics colleagues or people who just don't do this type of stuff and I show them an eBay page and they look at it and all of a sudden they start asking me, but how is it that this goes up and down and -- nobody even understands what's going on. And then I put up a scatter plot on the right, which just has, you know, the bids in five days. And then they start asking me, but why does it go down, why does it go up?

You can't even understand the rules of this type of auction without looking at something like this and understanding what the problems are. So, even just explaining or understanding rules, whether it's eBay or any other mechanism, or you're trying to move to a multi-unit, graphics make life a little more easy and transparent for all of us.

The second thing is we all have these regressions with -- sometimes they get very complicated. We're trying to build in these interaction terms and we get numbers for coefficients at the end. How do you interpret those? What we found, for instance, in the surplus paper was that using more advanced graphics, like looking at two-dimensional and using color and size and whatever, you're much more able to understand the real interactions of what's happening. So, for instance, we have the number of bids -- sorry, the surplus as a function of price and the shade being the number of bids or the number of bidders. So, you can see that you have different relationships and that's exactly what the interaction is telling us.

Just for the fraud point of view, this is a very interesting tool developed by our human computer interaction lab at Maryland, and if you know a map of the market on SmartMoney.com, this is where this came from. This looks at a bunch of -- about 10,000 eBay auctions on a bunch of different things and it's hierarchical. So, you have sports, within that you have golf, within it golf balls and you have lots of dimensions. You can use color, you can use size, ordering of these little bins, each little square, you can see is an auction. And what I used here for color is the seller rating. Green is high, black is kind of medium and red is really low. Where are the red guys concentrated? In the Rolex wristwatches, okay?

So, it's kind of interesting. Where are the very high-rated guys? In the Dell monitors. So, those are just types of visualizations very, very -- you see a lot of information.

The other interesting tool that we've put together is this tool that will couple the two pieces of information that you get from eBay auctions, which usually people just ignore one of them, the bid history and all the attributes. Everything is strongly coupled. You can filter, you can choose, you can look at things and everything is the same. So, you can explore not only the duration and the seller rating and whatever, but you can see the entire path of how the auction went, which brings me to my last point of dynamics. That is, looking at these curves, in a sense, and these are their derivatives, the first derivative and second derivative tell you about dynamics, how fast the price is going up and how fast it accelerates and decelerates.

Let me just put this up last. So, here's a bunch of about, you know, 150 auctions. These are the prices that we had and we segmented all these auctions to three groups and you can see that they have different bid paths, different prices going on during the auction. And more interestingly, when you look at the first derivative, then you're talking about dynamics. How fast was the price moving? Look, there must be more sniping going on in these auctions. They're all for the same PalmPilot in this case.

So, there's a lot of interesting stuff going on behind the scenes that we can capture by a different representation and we have papers -- if you Google my name and go to the website, you'll find a bunch of different papers on that topic, and I'll stop right here.

DR. ADAMS: Thank you, Galit. We're going to have Sean Peoples from Edmunds.

PANEL PRESENTATION BY SEAN PEOPLES

MR. PEOPLES: Thanks. First, thank you very much for having me and I do believe the first question out of everyone's mouth today has been, why are you here? Basically, from -- Edmunds, on a whole, basically, our mission in general is really to try to understand the dynamics of the automotive marketplace and provide kind of a view to consumers mainly. They're our main target audience. But basically it's kind of to the end of actually understanding the economics of the automotive marketplace.

We basically often kind of look into different analytics and different things, basically. We've got the Edmunds Price Index and true cost of incentives. I mean, just basically, we're always kind of looking at the effect of incentives and purchase prices in the marketplace. So, basically, we actually supply a lot of this stuff to businesses and analysts as well.

One of the areas that we actually have been looking into is generally data transparency and really it's kind of a fact on the purchase price of used vehicles mainly. You know, used car salesmen have a bad reputation for a reason. They've quite often used an information monopoly to -- you know, that's the kind of thwart a would-be purchaser, but in general, basically, that's what it does. It's really kind of (inaudible) theory from 1970, which was basically that because of that, people are actually hedging their bets.

So, on a whole, you know, that's on a transactional level, you may actually just be able to -- you know, to get one guy to pay more than what the vehicle's worth, but in the long run, basically, that's across all vehicle sales, you'll actually see kind of downward pressure because people are actually hedging their bets.

Really, there is kind of a significant lack of data transparency within the automotive marketplace and at large. The Internet's helped that a lot simply because of sites like Edmunds or even Kelley Blue Book or any of these other sites that you can go to, you know, that really try and provide a lot more information to consumers so they actually feel a little bit more empowered about information that may actually have been previously unavailable. I mean, you think even seven or eight years ago, invoice pricing was unheard of and now it's pretty ubiquitous.

But from a used vehicle value, I mean, there's three major items that actually affect used vehicles. Market values, basically those will just be that. I mean, there's a certain price that people are going to be willing to pay for a Ferrari versus a, you know, Ford Focus.

The condition that's actually becoming largely offset by vehicle history reports, the large number of photos that are available, even when you've got the separation of buyer and seller, in the case of an eBay auction, but equipment, basically, is kind of very quickly becoming a significant player in the value of the vehicle. Basically, that's because there's a significant increase in the complexity of how vehicles can be configured and, you know, it's the capability actually that separate out the buyer and the seller. So, there's a lot of information that actually has to be kind of taken on faith.

Some of the auction-specific conditions -- and I will actually admit that we have traditionally just been kind of researching offline wholesale auctions and comparing them to retail auctions and really just kind of an exercise for this particular conference have we looked into eBay in general. But, basically, the online consumer has vehicle history reports, you know, they've got valuation providers. There's probably 10 places you can go to see what the market value of a vehicle is. Most of the listings actually have a number of photos and a lot of them even actually that are on eBay have free Carfax reports and vehicle history reports.

Wholesale basically has kind of the same thing. You've got a dealer who walks in and he's got his, you know, book of used values, he's got an inspection report from an auctioneer, as well as, you know, they actually have the vehicle sitting right in front of them. But basically none of these tools really kind of confirms the equipment that's on the vehicle to the level of satisfaction of the actual -- of the buyer simply because there's too many variables. In an offline auction, there's just the speed by which the vehicle actually just comes through the auction. In the online environment, you basically just don't have physical access to the vehicle at all.

So, basically, we had previously looked at this topic to basically see if data opacity or the information monopoly basically just has caused equipment to depreciate faster than the vehicle value itself. And we actually just found that to be the case in the wholesale environment. Basically, kind of the inference from that would be that more complicated of vehicles would actually depreciate faster than less complicated vehicles.

So, given really the similarities between the online and the offline auction, basically, we should see the effect similarly across both of them.

This is really the previous research. Basically, what we looked at is that on an average the average transaction price of a vehicle was 25,500, the actual MSRP Was 27,400. This was a couple of years ago. And options, as a percentage, was 8 percent. Basically, the vehicle itself actually depreciates at a six-and-a-half percent difference than the actual vehicle itself. So, basically, we looked at it from -- there's basically $142 per vehicle left on the table at wholesale auction, which isn't a lot until you actually just realize the number of vehicles that are going through it and it's basically $3 million a year.

So, this time around, we basically just took a look at two diametrically opposed vehicles. Basically, it's the Ford F150 and the Honda Accord. Basically, the Ford F150 has 58 different trim variations and even within there, you basically just have a series of very high dollar options and many of them actually do have to be taken on faith. Things like axle ratios, basically, are completely unverifiable. I mean, that's even if you actually have physical access to the vehicle, you'd have to actually take apart the back of the truck to actually see if it's got it or not.

And then the Honda Accord basically just has three different trim variations. You've got leather interior and a navigation system, every single one of them that is available to view.

So, we took a look at the transactions over the course of a week and, basically, what we actually found was that the Honda Accord is pretty much in line with the industry average. We didn't actually have the average over the course of that week of all vehicles going across eBay, but basically the dispersion was pretty normal, but there was a very large disparity in the Ford F150s. It actually goes from significantly above retail average to being below the wholesale average. That's actually even further emphasized once you actually just look at the retained value.

So, basically, the Accord is three-and-a-half percent above the retained average. We looked at the transaction price of a 2001 vehicle and compared it to the transaction price that's observed in the retail, wholesale and eBay markets. And, basically, what we actually just found was that the F150 actually was pretty much in line, you know, half a percent above the retained value to being about 6 percent below in a wholesale market. So, it actually hurts the manufacturers and it actually hurts consumers in general, basically because you're paying a pretty high premium.

One of the interesting things, I'll back up here for one second, was that we did actually find or observe -- and it was a very small sample as compared to kind of a larger sample on the retail and the wholesale side, was that the eBay transactions were actually just tracking very, very closely to the wholesale transaction price. It's within 1 or 2 percent across the board. So, if you're buying a car, apparently it's a good deal.

Basically, the information imbalance or perception thereof, you know, really does exist, or at least we're actually seeing kind of evidence of it in people actually hedging their bets. It's the effects that you see in the wholesale auction and in eBay for that matter, are really just kind of symptoms of the larger issue. There's really no place that you can actually validate from a third party authoritative source to actually say this is how the vehicle actually left the factory.

Manufacturers actually have this data. Their concern is people actually that are doing competitive analysis on them, doing predictive analysis and basically, you know, just trying to figure out their market mix as compared to other people. There's a lot more kind of concern in that in automotive because it's a push market. But, basically, some of them are actually just kind of playing nice with us and actually just allowing us to kind of run some experiments with it.
But overall, basically, it's an area that there needs to be -- basically that's the equivalent of the vehicle history report in a vehicle equipment report that's validated from a third party.

Thank you very much.

DR. ADAMS: Thank you, Sean. I think you explained why you're here.

PANEL PRESENTATION BY JEFF HERRMANN

DR. HERRMANN: Hi, good afternoon, thank you. I'm Jeff Herrmann from Nielsen Media Research and I'm representing some Nielsen net ratings data. Today, we're looking at panel-based information and roughly I'll speak to two data sets for the panel. This is merely context data starting with some more general Internet-based information and then getting specifically into online retailing information. This data is more or less information to allow you to keep this in the context of the Internet as a whole, and I'll be curious for your feedback and discussions after the fact to see how some of these metrics would drive your consideration of looking at the server-based data, the actual transactions just from the site specific itself.

So, the first panel set will be looking at -- I'll just cut straight to the chase in the essence of time. This is a panel of RDD recruited panelists of approximately 40,000 people across the U.S. with a software tracking meter installed on their PC both at home and at work. As you can see, usage metrics-wise, the domains visited by a person from the home are roughly, on average, 61 websites visited, double at the workplace.

So, you have a significant increase in Internet usage activity in general at the workplace, and I wonder what the impact of that workplace activity is on bidding and want to know if anyone has incorporated that fixed location. The amount of attention is much greater on the Internet at work. That's your diversion at times or it's always there, it's always on, it's readily accessible. So, I was wondering if daytime bidding activity may see an up-tick given the metrics that we're seeing here at work.

Overall, just to benchmark, you have roughly 140 million people with access that were active in the digital media universe from home, 56 from work and in total, 150 million in the U.S.

Now, those are same trend data. This is all for you -- I'll leave this with you. You can have this information. Getting specifically to the local market distribution, you can see that this is just a ranker, once again, for context purposes of

Internet penetration by local market here in D.C. is actually the highest -- has the highest index of Internet penetration.

Demographics, demographics of Internet users primarily are representative of the entire population with basically a -- still a slight index here to more educated and upper incomes, but generally speaking, the online universe is representative of the total U.S. universe. This here is just a representation of the top 10 brands you see, and you can see that eBay has had a 14 percent increase year over year; Google, a 27 percent increase year over year; MSN, a 2 percent decrease. So, there is a redistribution of popular websites or popular brands, we call them, year on year, which is an interesting trend to see as the Internet changes the way that people actually behave with media. That's becoming more and more a part of their daily life.

So, roughly speaking to broadband, the most important point to consider about broadband and how broadband impacts usage is 80 percent of all webpage views -- so, of all the page views loaded throughout the course of time, 80 percent of all page views come from broadband subscribers. So, there still is an active dial-up universe. You can see it's declining. So, these are basically growth rates, and dial-up is going down, but the broadband connections are increasing but at a decreasing rate, of course.

Specifically to eBay, you can see of all online shoppers, eBay does drive a lot of traffic and has a significant presence in the online world. So, not only for straight retail, but eBay does actually have a significant presence. These are hard numbers to see, but you can see this is a distribution here of eBay's top-selling categories, August 2005, percent of purchases by category, percent of customers and percent of spending. Interesting to speak to automotive, you can see it's roughly about 3 percent of the total purchases on eBay are in the automotive category, but it drives 51.8 percent of the dollar volume on eBay. So, very few people spending a significant amount of money for eBay. It's an interesting area to watch.

And as far as looking at books and consumer electronics, the percent of purchases are lower than the actual percent of customers. So, you have that browsing effect of people just actually looking and may not actually go through and execute on a transaction, but they're checking here to see if it's available inexpensively. We have the ability to follow that. The good thing about the panel is we can follow the person's surfing pattern or their session through the life -- through the duration of that activity. So, we can see that maybe they were on eBay, didn't find what they were looking for and went on to Amazon or went on to another website to complete the transaction.

So, that's really just a simple overview that I had today and hopefully more to provoke thought about panel-based information versus the data that's been gathered here today.

DR. ADAMS: Great, thanks, Jeff. And last, but not least, Ana from the Bureau of Economic Analysis is going to inform us about something. I can't remember what.

PANEL PRESENTATION BY DR. ANA AIZCORBE

DR. AIZCORBE: Thanks very much, Chris, for a really fascinating day. Chris asked me to talk a little bit about how the issues that have been discussed today relate to the kinds of measurement problems that BEA cares about. The usual disclaimer applies. These are my ideas alone. BEA doesn't, in any way, endorse them.

I would like to focus on one particular type of measurement problem that requires information on how much consumers are willing to pay for goods, an area where the types of data sets discussed today may be able to shed some light. I have two charts that I would like to use to illustrate the problem. If anyone has ideas about how the data that you use in your world might be useful, it would be really great to hear about that later.

The issue is measuring quality change versus pure price change. It's an old problem. But what happens when you have heterogeneous consumers that have different ideas about what goods are worth? And what happens when these consumers are forward-looking or have an inter-temporal dimension to their problem? This is not only an interesting problem, it's an important one. It has potential implications for the speed at which price indexes for high-tech goods fall. So, let me just describe the problem briefly.

Let's say that we have two goods. One dies at time t and is replaced by another good at that very period. I've drawn their price contours downward sloping because that's how they look for PCs, high-tech goods, microchips and so on. It is important to remember that these price contours are tracking the same good over time, so the only time a quality change occurs in this simple example is when the new good comes in at time t.

Our challenge is to figure out the value and quality of the new good relative to the old good. If you think about this in a representative consumer way, it's not that difficult. You have a representative consumer that buys a little bit of every good every period, so every price that you observe in the market is a signal of how much the representative consumer is willing to pay for the good. The representative consumer's valuation of the differences in the qualities of the two goods is just the difference between those two prices.

But the problem is, what happens if people buying the two goods are different people? What if you're in a market for computers and you've got high-tech people and low-tech people? The person that buys goods two at time t is the high-tech person that's willing to pay a lot for the PC. The person that buys this one is the low-tech consumer that waits right until Dell's about to get rid of the computer and buys it at a very low price. If that's the case, then what this gap gives you is more like a valuation of the differences of their preferences rather than a valuation of the "quality difference" per se.

Conceptually, what you want in a market that has heterogeneous consumers is something like each person's individual valuation of the quality improvement in the new good, averaged over people. However, that's not what this gives you. So, there are basically two ways to look at this. You can look at this from a hedonic or BLP point of view, but that's not the relevant part for the kind of data that you have. The other way to look at it is what I'll refer to as a reservation price approach. Instead of thinking about a representative consumer that buys at least a little bit of every good in every period, let's think about the opposite extreme where you've got people only buying once. For example, the people that bought at time t are not the same people that bought at time t-1. It's sort of the analog to the new goods problem except it has to do with the entry of consumers into the market.

If you remember, in the standard case with the representative consumer and new goods, the numerical problem is that you don't observe a price for the good in the period before it was introduced. But there is a theoretical shadow value or reservation price that is consistent with driving demand to zero. The way to measure a new good in this perspective is to compare the difference between the price paid for the new good at time t versus the reservation price at t-1, and average that with everything else.

In our case, it's not only that there is a new good, but new consumers every period. And so, the analog for this case is that the people who bought at time t didn't buy at time t-1. They had a reservation price at time t-1 that was beneath the market price. The price index that we use utilizes market prices and shows fast declines. But if there are new people entering and exiting a market every period, the reservation price approach requires the use of those consumers' reservation prices, which we know to be less than the market price.

The scary thought is that the combination of heterogeneous consumers in an inter-temporal setting is that price indexes that use market prices will show faster price declines than what would be shown if you took these factors into account. This is an extremely important problem because quality improvements in IT goods was a major driver in the growth in GDP, over the nineties.

I have one paper where I looked at this in the context of Intel's MPU chips just to see how big this could possibly be. It turns out that you can use the same assumptions that Fisher and Griliches, and Griliches and Cockburn did when they looked at similar issues for prescription drugs. The reservation prices for goods with downward sloping price contours are bounded by the market prices under certain assumptions. Given those bounds, you can assume that the distribution of reservation prices between those bounds is uniform – as was done in Griliches-Cockburn – and then you can calculate an average reservation price for your price index as the average of the bounds.

If you do all that for microprocessors, the numerical differences are not small. A price index for Intel desktop chips over the nineties shows declines of about 30 percent per quarter. This is very fast. If you apply the Griliches-Cockburn strategy, the resulting price index falls 20 percent per quarter, on average. These are very large differences. But I hasten to add, this is just a first cut at the problem. We don't really have any hard evidence about where these reservation prices are. Finding the bounds requires assumptions that may not hold. And, nobody likes the idea of making this first-cut assumption that the reservation prices are uniformly distributed.

It seems that these data from internet auctions discussed today have the potential for providing insights into the determinants of reservation prices. Anything in these data sets that could tell us about consumer's willingness to pay for goods and how it varies across time could be useful for our measurement problem.

DR. ADAMS: Thank you, Ana. I call Ana's paper, did the nineties happen.

(Laughter.)

DR. ADAMS: I want to thank everybody and give everybody a big round of applause for a good job.

(Applause.)

(Whereupon, the conference was concluded.)