In the Matter of:

Economic Conference on Marketing and Consumer Protection

September 16, 2016
Final Version

Condensed Transcript with Word Index

For The Record, Inc.
(301) 870-8025 - www.ftrinc.net - (800) 921-5555

FEDERAL TRADE COMMISSION

ECONOMIC CONFERENCE ON

MARKETING AND CONSUMER PROTECTION

Friday, September 16, 2016
8:30 a.m.

Federal Trade Commission
Conference Rooms
400 Seventh Street, S.W.
Washington, D.C.

FEDERAL TRADE COMMISSION
I N D E X

FEDERAL TRADE COMMISSION
I N D E X (Continued)

FEDERAL TRADE COMMISSION
I N D E X (Continued)

WELCOME AND INTRODUCTION

(8:37 a.m.)

DR. JIN: Hi, good morning. Thank you so much for coming here. I know some of you have been at FTC before and some of you probably this is your first time to be here. Welcome you all.

I'm Ginger Jin. I'm the Director of FTC Bureau of Economics. When I took the director's role in this January, I had a strong feeling that FTC activity is very much related to marketing. Our bureau has over 80 Ph.D. economists, and we could benefit greatly from the marketing research community, the literature, the ongoing research in this community.

So I reached out to K. Sudhir and Avi Goldfarb just tentatively. To my pleasant surprise, both of them responded immediately and positively with many potential good ideas for getting together the FTC and the marketing research community. So I'm really glad that you can make today's conference. I hope will enjoy the conference and will find it interesting and be able to engage with us more in the future.

I would also like to thank all of you for responding enthusiastically to our call for papers.

We actually received over 50 submissions, which was really, really a good surprise to us. It also makes our scientific screening committee work really hard. Sudhir, Avi, Ganesh Iyer, and Andrew Stivers from FTC, they did a fantastic job putting together an agenda. But they wouldn't have be able to do so without your willingness to participate, to discuss, and to present the papers. So thank you all for doing that and being here.

I also want to thank INFORMS for cohosting today's conference, as well as Marketing Science Journal. I want to thank Laura Kmitch and Constance Herasingh. They probably are out of the room making sure that everything is running smoothly, as well as Stacy Awe. I think she is not here today, but she's from Yale and has been an assistant to Sudhir and very helpful throughout the planning of the conference.

A lot of my staff are on the ground here as early as 7:30. So I want to thank all of them. Ben Chartock, Jason Chen, Aaron Keller, Jennifer Snyder, Stephanie Aaron, Marilyn McNaughton, Maria Villofler, and Crystal Meadows. And they are -- we wouldn't be able to have the conference running so smoothly without them on the ground. So let's give a round of applause

to all of them.

(Applause.)

DR. JIN: And also thanks to the FTC admin team, event team, media team, for getting all the video and audio available for today.

So FTC has a history of over 100 years. It has a lot of interesting institutional features. To be honest, I didn't know all of that before I come to FTC. So I want to take this moment to just give you a brief review of exactly what we do at FTC, especially about marketing, about consumer protection.

So just to give you some sense, we know that FTC is in markets. Many markets would have one or more firms competing for consumers. So you probably have heard about competition and antitrust, which is one mission of FTC. I would argue that another even more important mission in the FTC is consumer protection. And that's because firms interact directly with consumers, and also the ultimate goal of preserving competition is to protect consumers.

Okay. And, moreover, if we think that firms -- if they feel like they are under unfair competition, they would have resources to go for private litigation and sort of seek some judgment from the court. It's really hard for individual consumers

to do so. Even if they can file class act litigation, it's going to be -- whatever the redress they can get will be shared with lawyers and some of them can be very aggressive.

So it's really important for federal agencies like the Federal Trade Commission to act on behalf of individual consumers to protect them from deceptive and unfairness, deceptive and unfair practice.

So FTC over 100 years actually has a lot of functions. The foremost is law enforcement. So I want to take this role. My daughter, who is 10 years old, asked me, Mom, what's your new job? I said, I'm going to be a policewoman. And she was saying how come you don't wear the police uniform?

So we're a law enforcement agency without uniform. We enforce over 70 laws against business practices that are anticompetitive, deceptive, or unfair. We can bring lawsuits in federal courts. We can also bring lawsuits in front of administrative law judge inside our commission. And if the decision of the judge was -- is contested, we can even hold commission hearings.

And after decisions from the court, we can enforce the final commission orders. It would also redress harm to consumers. So that's probably the majority of our work inside FTC. In addition, we also have rulemaking authorities. We can make rules for industry-wide practice. We also function as information collector. We watch out for new and problematic practices. We oftentimes, especially in the Bureau of Economics, engage in investigative research. We do a lot of research as well as policy advocacy.

So given that we are enforcing over 70 laws, it's probably very hard for me to give you a full list of all the laws we enforce. So I'll just give you a sub-sample so that you will have an idea of what we're enforcing.

We start from the 1914, the Federal Trade Commission Act, which gives us very broad jurisdiction over almost every industry, deceptive and unfairness and anticompetitive practice. And we enforce the Fair Packaging and Labeling Act together with FDA; and the Truth in Lending Act in 1968; the Motor Vehicle Information and Cost-Saving Act in 1972; and this is interesting, this Petroleum Marketing Practices is actually about franchisor and franchisee relationships in gas stations. So that was enacted in 1978.

And more recently we engaged in

Telemarketing and Consumer Fraud and Abuse Prevention Act; Children's Online Privacy Protection Act; Identity Theft Assumption and Deterrence Act; College Scholarship Fraud Prevention; Crime Against Charitable Americans; Do-Not-Call Registry legislation; unlawful internet gambling enforcement; U.S. Safe Web Act; Credit Card Accountability Responsibility and Disclosure Act; Patient Protection Affordable Care Act; and Restore Online Shoppers Confidence Act.

So of this history, it's just a subset of the laws that we enforce. You probably would get a sense that we actually enforce the laws in many, many industries, and recently more about online businesses in all kinds of actions.

So in terms of consumer protection, we sort of use two common legal standards here. In 1983, FTC actually published a clarification on Deception Policy Statement, which means we can go after the deceptions that are likely to mislead consumers acting reasonably in the circumstances to the consumer's detriment.

I will give you a few examples of what we mean by this legal language or unfairness. In 1980, we clarified that it's going to be a three-prong exercise. It has to generate substantial injury or likely to generate substantial injury, not reasonably

avoidable by consumers, and not overweighed by countervailing benefits to consumers or competition.

So let me give you a few examples of exactly what we do so that you will have a sense of the activities here. So I will first go over some examples and conclude with some challenges we face today. And I hope you can help us addressing those challenges.

So the first example is fraud. The Bureau of Economics actually worked with the Bureau of Consumer Protection to conduct nationwide fraud surveys for three rounds. And actually the fourth round is ongoing right now. So I'm -- here I list a few reports from those surveys.

In the latest one that we have data on, which is 2011, we actually observed about 10.8 percent of U.S. adults or 25.6 million people were fraud victims. This is -- I don't know whether you think this is a big number or small number. It was kind of shocking and a surprise to me when I read the number, and in total we estimate there are about 37.8 million incidences of fraud during the year of 2011.

So our fraud survey also gives us some sense about what type of fraud are most popular on the ground. Okay? So this graph is sort of showing you

the top, I guess, 15 to 20 types of fraud by number of victims. The number one actually is also number one in the last round of fraud survey. It's weight loss products. Okay? And it follows by a lot of creative scams such as a prize, promotions, bidding buyers club, internet services, work-at-home programs, credit repair, and on and on.

So you can see that the frauds we are supposed to police are really widespread and can take many forms. So the challenge we face is how can we attack those frauds given there are so many going on in the market and how can we educate consumers to avoid those scams, and eventually how to penalize and deter those scammers, especially when they are fly by night. Okay? They sort of gather all the money and already spend it by the time that we catch them. How can we really penalize them and deter them is a quite important legal as well as economic question.

So that's the first example. The second example is actually a lot of cases that we have brought before were about deceptive advertising. Some of you may have heard about this or even own a car from Volkswagen. Okay? So this is a typical ad that Volkswagen put up for their really clean diesel. Okay? It turns out that it's only really clean when it's tested because they have a defeat device which is a software hidden in the car, and that software would understand that the car is under a test mode and sort of trigger cleaning process inside the car. But when it's not in the testing environment, it actually can generate NOx as much as 40 times above the federal standard, which would have a significant consequence for the environment as well as for people's health.

So in March of this year, FTC sued Volkswagen over deceptive diesel claims. And thanks to our collaboration with a lot of other federal agencies such as DOJ and EPA, were able to reach a historical settlement which is as much as $10 billion, which means they -- Volkswagen is willing to pay up to $10 billion to consumers who have been deceived by these ads, and they all have options to sell the car back to Volkswagen with substantial monetary compensation, or they can have the car repaired and still own the car. And even with that option, they will receive significant monetary compensation.

So this is quite a victory for FTC and eventually to all the consumers in the market. So that's the example of deceptive advertising.

Another example is privacy protection. If you had been here yesterday for the Disclosure

Workshop, there has been a lot of attention on privacy, consumer privacy and how to protect it. So a recent case we brought is for Practice Fusion, which is a cloud-based electronic record management company. So we allege this company started to collect patient evaluation of doctors since April 2012. For over a year, the website has collected a lot of consumer reviews. In April 2013, they decided to go live with over 613,000 consumer reviews.

However, some of them include highly sensitive personal and health information. And at the time that they entered those reviews, the privacy notice they received did not indicate there will be a public display of consumer reviews. So we think this has violated consumers' privacy, and we are able to reach a settlement in June of this year with a 20-year order to constrain this company.

The fourth example is online endorsement. I know many of you have done very interesting research about online activities, online advertising, online endorsements. This is an area that we are very active in watching and policing.

So one example is a case we brought in 2014. We alleged that Sony Computer Entertainment America, which is a branch of Sony Company, falsely claimed

their Playstation Vita console has game-changing technology features. Okay? It turns out to be false.

In addition, their ad agency -- I think in Los Angeles, Deutsch LA, misled consumers by urging its employees to create awareness and excitement about this console on Twitter without disclosure of their connection.

Okay. So we think this is not acceptable. So we reached a settlement with them. For both companies, we have cease-and-desist orders, and Sony also agreed to pay either $25 in cash or $50 in merchant credit to buyers that have bought this console before June of 2012.

The last example I want to give is about multi-level marketing. Okay? I don't know how much you know about multi-level marketing. It's turned out to be a very big industry. So this year we brought a case against Herbalife, which is the third largest multi-level marketing company in the world. We allege that they deceived consumers into believing substantial income from the multi-level marketing business opportunity, which is a deception count.

We also allege that they incentivized distributors to buy products and to recruit others to join and buy products so that they can advance in the

company's marketing program rather than in response to the actual consumer demand. And this is an unfairness count.

So we're able to reach a settlement in this July. After a two-year investigation, that settlement included a $200 million payment from Herbalife for consumer redress as well as restructure its business from top to bottom. We hope this is a historical case that will help to shape the whole industry of multi-level marketing.

So with all these examples, you can see that we cover a lot of areas. We try to keep up with the business practices going on in the market. We also face a lot of ongoing challenges. The first one is how to detect potential violators. We sort of have some sense -- we have a lot of experience in dealing with offline violators, but our knowledge is that many of them have moved to online with probably more decentralized networks, with more creative actions in desktop, in mobile environments. And so really we want to engage you in understanding the marketplace and try to think more about how can we do a better job detecting potential violators for both online and offline markets.

Another question is how can we link consumer misperception to firm behavior. In many cases, we observe outcome. Those outcomes could be driven by many factors, including the firm's wrongdoing, but as well as other noises in the market. How can we distinguish all those things and use the information we have to go after real violators? How can we define the measure of consumer harm and countervailing benefits? And that is already hard in the offline markets, but it's become even more challenging in a world of big data and connected things. So we really want to hear your research in this area.

There also is a sort of policy question if we are sure that there's something wrong going on in the field, we want to change the market. Should we discipline the firm? Should we educate consumers? Should we do both of them, given our limited resources? So that's probably a more policy-oriented question, but it's also very related to our understanding of the market and about the potential effect of our policies in this area. And how to regulate a market when consumer knowledge and business practices are both evolving.

And we know that consumers care about privacy from many consumer surveys, but they also behave in a way that seems sometimes inconsistent with

what they say in those surveys. So consumer knowledge definitely is evolving and businesses probably are adjusting their practices to this kind of consumer demand. With others evolving and probably both of them will change given our position in policymaking. So this is a very dynamic and ongoing environment. It's really challenging for us to think about the interactions between the different players here. So, again, that's -- we really want to engage your thoughts and your creative thinking on exactly how to address that question.

And, finally, if you have any ideas or any comments or suggestions as to how we can better engage with your community, how can we learn from your research community, it's really, really important for us to keep up with that literature.

So with that, I will mention a few sort of logistical things that we have to say, and then we'll move on to the real content of papers.

You probably already have the pamphlet about wi-fi and information. Okay? And this is a federal building, so if you are going to leave the building unfortunately you probably have to go through the security again. Okay?

And this conference will be recorded by our

technician, Jennifer. So we would require everyone to speak into the microphone, the presenter will speak into the microphone. We also will have walking mics around the room. So when you want to ask a question or want to make a comment, we hope that you can speak to -- can wait for the microphone to come to you and speak to the microphone so that we can record the whole conference.

And there are actually restrooms on this floor. If you go out of this conference room and past the glass doors you just used to come into this floor, there will be a restroom on your right-hand side.

In case of emergency, if emergency occurs and requires you to leave the conference center but remain the building, please follow the instructions provided over the building's PA system. If an emergency occurs that requires the evacuation from this building, alarm will sound and everyone will have to leave immediately upon the alarm. And we are supposed to leave in an orderly manner, not rushing to the door in a congested way.

And so in that case -- and hopefully it's not going to happen, but in that case, we'll need to leave the building through the main 7th Street exit. After leaving the building, we'll turn left and

proceed south past E Street and there will be a FTC emergency assembly place. Okay? And so we'll be -- we'll remain there until further instruction. If you notice any suspicious activity, please alert the building security.

So, finally, please be advised that this event will be recorded and we'll have a transcript later on available on the website. What you provided to us might be -- the event might be photographed, webcast or recorded, and what you say here, your image, and what you submitted here will all be subject to potential posting on FTC.gov or at any social media website related to FTC.

So with that disclosure, the privacy notice, I want to thank you all for coming here. So we'll kick off with our first paper. Our plan is to have 40 minutes per paper. So we'll have 25 minutes for presentation and 10 minutes for discussant, and hopefully we'll have five minutes for floor discussion.

So our first paper will be presented by Garrett Johnson from the University of Rochester about the impact of Privacy Policy on the Auction Market for Online Display Advertising. So that title basically captures a lot of key words I just said about the FTC's business. So that's a great start.

Garrett?

(Applause.)

SESSION ONE:
THE IMPACT OF PRIVACY POLICY ON THE AUCTION MARKET
FOR ONLINE DISPLAY ADVERTISING

DR. JOHNSON: All right. Well, good morning. Very excited to have this conference. I really want to thank the organizers for doing this. This just makes me so happy to see the logos of Marketing Science and the FTC close together.

This is an area -- the area of online display advertising that I think is really ripe for research, but also has a lot of challenges from a regulatory sense and from a researcher's sense. Firms in this industry have essentially rewritten the social contract as it pertains to online privacy with a lot of kind of benefits and maybe some harms that have accrued from that.

To get things started, usually I would start with my own laptop and I would go to The Chicago Tribune and I would open this add-on to Chrome called Disconnect that shows you all the different firms that know that I visited The Chicago Tribune. And you'd see that there'd be about eight companies or 10 companies that would know that. And then I would press the unblock tracks button and the amount of companies would spread like amoebas in a Petrie dish.

And after a few seconds, you'd have about 100 companies that know that you visited The Chicago Tribune.

Now, this helps make salient, you know, we talk about tracking, but it really helps make salient just the amount of tracking that happens online. And as you spend more time on the internet, you become intertwined in this web of companies that are in some cases just kind of benignly tracking for the purposes of ad measurement, and in some cases benignly tracking for the purposes of measuring traffic online. But in other cases they're really trying to create a rich profile of who you are as a consumer, what your interests are, so that that information can be used to enrich ad targeting.

All right. So start with an overview. So as you know and probably as the reason I'm here, U.S. regulators are interested in possibly regulating this industry. And that is at all levels of Government from the White House to the FTC to -- there's a bill in the Senate, there's a bill in the House of Representatives. So all levels are interested in this topic.

And it's a really challenging topic because on the one hand you have the privacy concerns.

Certainly some users are very concerned about privacy practices that are prevailing in the industry. And on the other hand you have an industry that is very dynamic, that has grown from $1.7 billion in 2002 to $7.9 billion in 2013, which is kind of when this paper was written. Nowadays it would be about $11 billion. So for FTC people, that's about the same order of magnitude as the Volkswagen settlement.

So the goal today is to measure the effect that privacy policy would have on advertiser and publisher profits. So you'll notice that there's one thing absent from a welfare calculation there, and that is the welfare of consumers. That's a really challenging question to also tackle. I have some followup work that I intend to spend a couple minutes at the end of the presentation talking about that looks at the consumer side. But really the focus of this paper is to quantify the effect on the firm side.

So the way I'm going to do this is I'm going to use an empirical auction setting and I've got to gather a lot of data from realtime bidding and this advertising auction marketplaces, and then I'm going to use a structural toolkit to construct a world with privacy legislation.

So I think one thing that I'm excited about for this paper is that there's been a lot of growth in the economics and privacy literature, but this paper, to my knowledge, is the first paper to take a structural approach to answering this question. And I think it's actually a really natural set of tools to use for privacy legislation because usually these sort of regulations are irreversible.

And we would like to try to be able to construct a world ahead of time that would inform what we think would be the consequences in such a privacy environment, or such a policy environment, and so I think the structural toolkit is going to be very helpful in this regard.

Now, the challenge, of course -- and there's a myriad of challenges in this project -- the main challenge is that I don't get to observe the information that firms have about users. And so I'm going to model that as unobserved heterogeneity in the marketplace and I'm going to extend models of unobserved heterogeneity in auctions to be able to answer the question.

Now, the high-level results is that I -- my model shows that the surplus in the industry would fall by something on the order of 40 percent. Now, I'm someone that's very motivated by these policy

questions. It's really important for me to get these numbers right. I think that, you know, to be very transparent I've got more to do to show -- to advance those numbers to really -- to really nail them, and to more importantly show how those numbers can vary under different scenarios. So I think there's more work to be done there, but it gets the conversation started.

All right. So just -- because I don't have a lot of time, I'm going to move fairly quickly through things. So there's a number of papers that have looked at privacy policy, that have looked at the online display marketplace. The intersection, there's fewer papers.

One notable paper is by Avi and Catherine that looked at a switch in the European marketplace where advertisers were suggested that they shouldn't be tracking. And Avi and Catherine found that that decreased ad effectiveness on the order of 60 percent according to causal effect marketing surveys. So that was a really helpful way of starting the discussion off. But my paper is going to take a structural approach to try and quantify this effect in dollars and cents.

All right. So I want to begin by giving you a taste of what the industry looks like. Part of

what's exciting here is that the industry has really changed a lot from the old days of buying and selling advertising. So on the one hand you have users who are people like you and I that are creating the opportunity for ads to be sold.

And in this marketplace, the unit of advertising is an ad impression, which is really fine-grained. It's a single ad on a single computer for a single user on one position for one page load. So any time you're loading the page you're creating more ads, and the ads that I see are a different marketplace than the ads that you see.

Now, on the sale side of the marketplace you have publishers like The Chicago Tribune and The New York Times; and the buy side, of course, you have the advertisers. And they're going to meet in some marketplace in the middle. Now, that marketplace predominantly takes two forms. One is the guaranteed contract marketplace, and the other are ad exchanges. On the guaranteed side, the sort of contracts you'd see are basically bulk buys of advertising ahead of time. So a contract that you would see would be Coca-Cola contracting with Yahoo! to purchase every user that visits the Yahoo! front page in the United States on a certain day. And that would come with a price

tag, of course.

Now, the thing with contracts is that they have contracting costs, and the contracting costs can be really high in this marketplace. So a second approach is to use a realtime auction hosted by a set of firms called ad exchanges. And that is where things have really changed in this marketplace from the sort of handshake deals to basically computer-mediated commerce that determines how ads are bought and sold.

Now, my data set comes from an ad exchange, and really that's where the tracking happens. Right? If you're trying to find people that visited Madden Football in the past, you're not going to buy like a bulk buy on Yahoo! What you're going to want to do is try to find these people across all webpages on the internet and you're going to buy them on the ad exchanges.

All right. So to participate in the ad exchanges, there's two main ways of doing it. There's the one way which is realtime bidding. The second is what I call offline bidding.

Now, what the realtime bidders do is they're going to evaluate and bid on the individual ad impressions that are coming down the pipes. And to do

29

so, they're going to employ computer algorithms. And they need to do so because this marketplace clears in less than .01 seconds. And so we can't hire, you know, undergrads or MBAs with fast fingers. We really need computers to be cranking through this data.

So this is the prominent way that people buy and sell ads now in these marketplaces. My data set is about five or six years old, and so much more of the data is from offline bidders. Now, what they do is to solve the speed problem, they basically operate as proxy bidders that specify their bids ahead of time. So they're going to specify rules like the target audience that they're going after, the fixed bid that they're going to submit again and again, and then they're going to submit a budget. And the way that was operationalized at the time is they would just randomly submit their bid over time to spread their budget across time.

Now, the important thing to realize is that both these bidders are going to employ user tracking information, and the bid data is going to look very different. On the realtime bids, you're going to see basically continuously distributed bids whereas an offline bidder you're going to see the same bid again and again. And so the challenge is to model how these

30

two different kinds of agents are using information.

All right. So let me talk to you a bit about this identification. So, big picture, unobserved auction heterogeneity refers to the case where bidders know more about the object for sale than we do as the modeler or as the econometrician, and in this case we'd have some observed heterogeneity. So I get to observe that ads are being sold on certain publisher sites and certain ad slots, and I get to observe a little bit of information about users like the country of origin. But the unobserved auction heterogeneity in this case is the tracking reports that advertisers have about users.

Now, the problem is that the existing models of unobserved heterogeneity, you can just sort of think conceptually it's going to be pretty hard to kind of find what's invisible in this marketplace. And so the existing models require that there is no binding reserved priced and that you observe all the bids.

My data is really unpleasant in that regard because 80 percent of the time the reserve price bid binds part of me. One percent of the time I only observe a single bid, and 10 percent of the time I observe at most two bids. So really from a

31

perspective of modeling it's kind of like tying both hands behind my back.

Now, what helps me is that I get to see the same users being bid upon again and again and again, and using that panel structure it's going to allow me to try to disentangle what could be coming from the panel -- sorry, from the tracking reports.

All right. So let's start with the offline bidders. Just to remind you what they're doing is they're specifying a target audience that you can visualize. There's a space of users and the red circle is the space of users that the advertiser cares about. And they're going to be submitting this fixed bid with a certain probability.

Now, the way I conceptualize this exercise is that it's really important to know the size of this target audience. And the reason for that is that you can imagine that right now the advertisers have a lot of information, including gender. And so you can think that, let's say, Gillette is bidding on men, they're bidding $1 for men and men occur with a half probability in the population, I'm reliably informed. And so what I'm going to do in the counterfactual is I'm going to say that the bid is going to scale down by the size of that audience. And so in the

32

counterfactual the advertiser will be bidding 50 cents.

So the crucial thing is that this bid is going to be scaling up or down based on the size of this target audience, and so I want to quantify that. There's going to be some challenges, though. The first challenge is that if these people are randomly submitting bids, then I'm only going to observe a subset of users that are hit with these bids.

The second challenge is that this is a competitive marketplace where I observe at most one or two bids. So there will be cases where there's competition that sensors the highest bid, and so I don't get to observe users for which I is interested.

So the way that I solve this problem is I basically say, well, this can be -- this world can be understood to contain four types of users, people that nobody wants, people that only Advertiser I wants, people that have some overlap between I and I's competition, and those for which I is uniquely -- sorry, the competition is uniquely interested.

And so I'm going to treat this as a mixture model, and I'm going to identify this using repeated observations. The basic intuition is that if I see the same user again and again and they only get reached by Advertiser I, then it becomes increasingly likely that they're only in I's set. And so these repeated observations are going to allow me to pin down the size of these different elements in this figure.

There's a question from -- yeah?

AUDIENCE: So the history is attached to the eyeball? Everybody gets, like, their trackings?

MR. JOHNSON: So, in this case -- in this case I'm going to be able to -- the nice thing about this is that this model allows for a fully general overlap between I and I's competition. And so this model can accommodate cases where advertisers are potentially getting different information, number one, and, number two, if they're interested in different characteristics of the users.

AUDIENCE: And in reality -- no, I'm just wondering, like, is it sort of, like, okay, here's an eyeball and the -- does the -- how do I know what that eyeball -- I'm sorry. So the reality if there's an eyeball, do they have information? Is it different information? Is it the same information? Does the auctioneer offer the same information?

MR. JOHNSON: Yes. So, in reality these ad exchanges get some information -- they typically have

some information that they can make available to everyone, but oftentimes advertisers bring their own information to the marketplace. So a specific example of this would be retargeting. So you look at a pair of socks on Macy's and Macy's hunts you down for the rest of your real life to convince you to buy a pair of socks on the internet. Other advertisers don't have that informations but Macy's has that information.

So that's -- you know, that's actually another challenge in this setting, is -- you need to make some simplifying assumptions about who's got what information.

AUDIENCE: Just to clarify, it's not Macy's that has -- sorry. It's not Macy's that has information about the ad network, right?

MR. JOHNSON: We're sort of splitting hairs here, but Macy's or Macy's ad agency or somebody somewhere who's representing Macy's has got that information.

All right. So very good questions. And, you know, one thing you're probably realizing if you're new to this area, there's a lot of nuanced stuff going on in the institutions that have really changed a lot in the last five or six years. So

that's, you know, obviously a challenge for researchers and certainly a challenge for regulators as well. Heck, it's even a challenge for industry.

All right. So I've talked about one type of bidders, which are these offline bidders, and I've told you that one nice feature of the model is I'm able to have a lot of richness to advertisers targeting different users.

In the realtime side, I have very rich bidding space, but I'm going to have to pin down some of the common, unobserved heterogeneity a bit more. So what I'm going to say is that in the realtime space the valuations of the advertisers of the product of two terms. The X term is an idiosyncratic term, which varies continuously, and then there's going to be the second term which is an unobserved heterogeneity term.

Now, the assumption that I need to make in this case is that that unobserved heterogeneity term is fixed for a given user. And so you can think of this as capturing to some extent a user's responsiveness to advertising and their income that they have to spend on various things. But because I'm making this assumption, it's going to not fit very well with the world where BMW is going after rich people and Kraft Mac & Cheese is going after poor

people. But, again, given the data that I have I think this is as rich as I can make the model.

So under some assumptions I can pin down these two things, and the counterfactual I'm going to run is I'm going to shut down the variance in the unobserved heterogeneity component, which models the tracking reports, and I'm just going to allow the idiosyncratic term to vary.

All right. So let me explain how I go about identifying this model. You know, really what I need to do is kind of identify this by a certain amount of brut force, because it is so challenging to disentangle this. So what I'm going to do is -- you ever run a thought experiment where if you observe the same user again and again and again, like hundreds of times or thousands of times, then for that same user you're holding fixed the Y component. You're holding fixed their unobserved tracking component.

But so that's going to tell you -- the variance in the bids is going to inform you as the X component. However, if I hold -- if I look across people and I look at some quantile like the maximum, I can sort of sort everyone in the audience by the maximum bid that they achieve after observing 1,000 bids, let's say, and that's going to tell me something about the unobserved heterogeneity component. So that's how I go about identifying the model.

Now, there's a question over there?

AUDIENCE: Yeah. Is your counterfactual -- how does that correspond to an ad blocker? Are you basically implementing ad blocking into this counterfactual?

MR. JOHNSON: No. Because ad blocking -- so this -- you know, our whole story of, you know, an auction runs in .01 seconds, the story is really boring with ad blocking. It basically stops when you install the ad blocker. There's no auction, there's nothing that happens. And so -- unless ad blocker starts selling ads, which apparently they want to do. Yeah, they're a delightful company.

So, yes, the ad blocking basically, the answer is -- we know the answer is zero until maybe yesterday. The answer becomes something for ad blocker, Ad Block Plus.

All right. So I'm not sure how I'm doing on time, but someone will yell at me eventually. So let me tell you a little bit about the theory. It just kind of shows you what's going on in the background.

So in a structural auction paradigm, what we do is we observe a bunch of bids and we want to

transform those bids into the valuation of advertisers or of the bidders in the auction. And then in the counterfactual we're going to -- we're going to make some change to the environment holding the valuations fixed and then simulate our toy model of the world.

So what this assumes is that we have some model that connects valuations and bids. And here the auctioneer uses a really unusual mechanism in that it makes offline bidders play by second price bids, which means that they're paying the second highest bid or the binding reserve price, and it makes the realtime bidders play by first price rules.

Now, why are they doing this? I don't know. It's a little crazy. Most of the industry now uses second price rules. By the time -- I'm just speculating that maybe they're trying to penalize these more agile first price bidders a bit.

So kind of the simple version of the theory is that it's a dominant strategy to bid your own valuation for the second price bidders. The first price bidders want to shade their evaluation so they maintain some surplus. The challenge in this setting is that if you are a first price bidder and you're facing someone that's bidding a dollar again and again and again, you never want to bid .99 centers because

you should bid $1.01 and the chance that you win improves discontinuously. This means that there's some bids that theory tells us that we should never observe.

To just kind of visualize that in the simple world where you've got two uniformly distributed first price bidders, the optimal strategy is to bid half your evaluation. Now, if you then put in some person that's bidding 25 cents half the time, then at a certain point you cross an indifference threshold and the optimal bids kick up and you observe this gap where your theory tells you you should never see bids.

Now, the challenge is when you work with real-life data is that these are pretty small stakes auctions, it's pretty costly for advertisers to learn, so of course these people go and they bid in these gaps. And so a big part of the headache that's kind of held up this project is to think of an intelligent way to model this to gain as much information as possible and to be able to do so reliably.

All right. So the results then, as I said at the beginning of the presentation, it's something on the order of 40 percent. It decreased if you were to ban tracking outright. It's felt a little bit more by the advertisers and the publishers, though pretty

evenly.

The scope of the results is that I'm focusing on the top three websites in the data, which is about half the revenue. So to kind of a do a back-of-the-envelope calculation, at least in 2013 you would take $6.8 billion, multiply it by the 20 percent share that does this realtime auction stuff, multiply it by the impact on the industry, and it's something like a half a billion dollars.

Now, today the revenues are closer to $11 billion. The auction share is closer to 40 percent. And so you're looking at $1.5 or $2 billion impact on the industry.

All right. So since I have a little bit of time, I wanted to tell you about some followup work that I'm working on that that I think will be interesting to this audience and that I hope to talk a bit more offline.

So in this industry, the industry tried to do -- well, it did do a self-regulation mechanism. You may recognize this little triangle with the I in the middle from the top corners of the banner ads that you see on the internet, those of you that aren't blocking ads. If you click on one of those things, it's going to take you to a website that's going to tell you about the benefits of personalized advertising, but then will allow you to opt out of this form of advertising.

And I was able to obtain a proprietary data set from an ad exchange to take a look at this question that's quite recent. This is a year ago. And I think this is important to look at because as Ginger remarked, when people -- we asked people about how much they care about privacy; everybody says that they care a lot about it.

And when you look at if they take any sort of action that is consistent with those beliefs, you realize that a very small minority does that. And I think that, you know, both research perspectives have something to teach us, but certainly from a regulatory perspective what you care about is what's actually going to happen in real-life. And so I think that discussion should be informed by this revealed preference study.

So the big questions I answer here that maybe -- or that I'm trying to ask here but maybe I'm not going to tell you with the stenographer writing, is, first of all, how many opt out. It's actually a very, very few. I've tried to get a sense of who are these people that opt out. I looked at how the

marketplace outcome is different for those that opt out, and it's about -- pretty comparable actually to what I'm seeing in this project. And then I look for heterogenous impacts, which I think is really important from a regulatory perspective. And it turns out that certain types of publishers have a lot more of these users than others. So, again, I hope to talk more about this project offline.

But getting back to the main study here, the goal was to try to enrich a policy discussion that I think is very interesting and very important with some numbers that try to estimate the impact of this policy on the industry.

Now, again, this paper is the first paper to take the structural tools to a privacy policy question. I think it's a set of tools that can be very helpful to answering these questions. The takeaway from marketers is that there's just such an exciting change in the industry from measurement to the way that ads are bought and sold, to the privacy questions. Advertising has really pushed the frontier of what is possible in the last decade. But to use a Star Trek analogy, the frontier is starting to push back with people blocking ads, among other things. And so that creates challenges that are not just

shared by regulators and us as academics, but also shared by people in industry.

So, again, I thank you for putting together this conference which I think speaks to very important issues and it's really exciting as someone that thinks of these issues as my bread and butter to see the leaders of the field pushing the same research questions. So thank you. I look forward to the discussion and some questions afterwards.

(Applause.)

DR. JIN: Thank you, Garrett. That's an exciting research agenda. It's extremely relevant for FTC. So our discussant is Doug Smith from the Bureau of Economics at FTC.

DR. SMITH: All right. So I'm just going to take a minute here.

DR. JIN: While Doug is pulling up his slides, I want to remind everyone that we want to record everybody's conversation here. So please wait for the mic to come to you. If you have a question, please raise your hand, we'll try to come to you immediately. We also have flash cards at the back for presenters and discussant so that you will be tracked by time. Thank you.

DR. SMITH: So, hi. So I'm discussing this

paper that Garrett just presented very well. And I don't have a lot of time. So I just, you know, first wanted to say that, you know, this is a very clever approach to dealing with this problem. Garrett has drawn from a lot of different auction literature to kind of deal with the very specific market he's looking at, and the way that the pieces fit together to identify these values is very impressive.

You know, basically all this machinery is really for the purpose of just when an advertiser is bidding, what is the actual value that they assign to this potential ad? And so one thing I wanted to highlight about the nature of the exercise that he does is that when he's looking at the counterfactual, one thing that there just isn't data on is what the advertisers would value -- how much the advertiser would value a user who they aren't targeting.

So they observe all this bidding on certain people, but the counterfactual has to deal with the fact that, you know, without knowing who's who, you're going to be buying sort of an ad with an expected value that covers just sort of the average cross population. And so knowing what the person -- what an advertiser would value somebody that they're not generally targeting is sort of crucial to figuring out these values.

And the paper -- you know, the paper -- Garrett talks about this in the paper and he knows that the value could really be anything from zero to the reservation price. And I think kind of as a sort of, you know, very clean way to do it, he does the counterfactuals estimating that the value for sending an ad to a user who you're not targeting is just zero.

I think, though, that this really is an area of uncertainty. This data isn't really telling us anything about this. And so in that sense, you know, a useful exercise would probably be to provide estimates using a reservation price or something sort of analogous to that as an alternative and just sort of seeing how much matters.

So you can imagine they could get very similar results, in which case we know this uncertainty doesn't affect much, or potentially get something slightly different and then know that there's sort of a dimension that we don't understand.

So besides just that comment about the paper, I want to sort of step back a little and think about what this tells us about the policy question here. And, you know, as Garrett mentioned, you know, obviously advertising publishes just part of the

picture. We need to also understand the effect on consumers. And generally, you know, people talk about these things as sort of two components to -- particularly as an economic question. You know, is the tracking here, is it making the pie bigger, you know, so that everyone can benefit? And this would be basically through better matching, you know, more or better matches.

Alternatively, or perhaps as well, is targeting and allowing companies to take a bigger portion of the surplus generated through -- generally through price discrimination. And then, you know, you also need calculations maybe account for consumers privacy, just specific preferences.

But, you know, so this is sort of what Garrett has done is an input into this process, but there's these other components to consider as well. Oh, I'm sorry, I want to say one other thing about this. But I think that, you know, looking at data, particularly how firms end up making transactions and what prices and such we can probably actually get some interesting insights into these questions but more on the firm level data. But I think this is something that I'd encourage people to sort of start thinking about how to explore.

So another aspect to sort of think about the bigger picture is this question of how are consumers actually going to react to these privacy policies. So one thing I think that maybe is worth explaining a little bit is sort of what the policies are that all under consideration.

So Garrett actually considers three sort of alternative policies to the status quo. One of them is just to allow consumers to opt out from targeting. And he draws from various sources to sort of get a ballpark of about 10 percent of consumers he predicts would opt out. Another possible policy is just an opt-in policy where, you know, unless you say you're willing to be tracked, you won't be tracked. And, again, using various studies, he sort of estimates maybe around 90 percent might decide not to opt in. And then the third policy consideration is just a blanket prohibition, which would be, you know, a default by automatically 100 percent not in.

So -- but a thing that -- you know, and this is, again, something that Garrett raises in the paper, you know, companies may adjust the incentives they offer for people if they, in fact, face a significant number of untracked customers. And so, you know, you can imagine companies sort of trying to get you to opt

in. And I think that that's an area where it really needs to be explored further and provides some interesting potential for future research opportunities.

Okay. That's actually basically all I had to say. You know, I think, again, this is a really interesting contribution both methodologically and sort of helping us start thinking about this policy area. And something that I didn't realize Garrett was going to mention repeatedly, but he has a really good point, is just that these things are evolving so much. And so it will be very interesting to see how in similar exercises what kind of answers they get in the future. Thank you.

(Applause.)

DR. JIN: Thank you. We still have time to pick up questions. If possible, I will ask you to state your name and affiliation first and then ask questions. Thank you.

DR. LIAUKONYTE: My name is Jura Liaukonyte, I'm from Cornell. So one of the things that I've learned that was amazingly surprising about realtime bidding is how much ad fraud there is. There's a ridiculous amount, like 50, 40 percent, which is essentially publishers stating they are putting ads but they are not putting ads. How does that affect the welfare calculations, if at all?

My thinking from sort of equilibrium perspective is that if there was no ad fraud, the prices would be higher. Right? So the advertisers are already incorporating that information in their bids.

DR. JOHNSON: Yeah, I think you're right that there's -- if you expect -- if you expect the value of an ad to be a dollar as an advertiser and then you expect the ad to be kind of true half the time, then you're going to deflate your bids accordingly. So hopefully they're accounting for that in the marketplace. But the lack of transparency makes that really difficult. I've heard that the numbers aren't quite so high.

You know, one more thing that's changed is that now the marketplace allows for cost per viewable impression payment models rather than cost per impression payment models. So that mitigates those concerns a little bit. Yeah, it is kind of the wild west even for industry people in this marketplace, especially if you're getting away from kind of the big three, the Facebooks and Googles and whatever, Yahoos of the world.

You had a question?

AUDIENCE: Okay. So this gives us a sense of how much a certain kind of policy might affect industry. And you're leaving consumers aside. Given this, you know, is there anything you think you can say about consumers in terms of the model or in terms of at least how bad things would have to be for consumers in order to make the policy worthwhile?

DR. JOHNSON: Wow. So the consumer side -- so one way you can think about this is, you know, this industry is about $11 billion. There's -- I'm Canadian -- there's 350 million-ish people in the United States. So we're looking at like $30 or so per person, right? So that kind of brings up the point that you made, which is, you know, the -- you might think that the firms could find some way of rewarding consumers.

It's pretty hard to in our current kind of financial infrastructure to do micro-awards commensurate with, like, $30. You know, if that infrastructure changes in the future with digital currency, you know, we're into unchartered territory where perhaps advertisers could try to compensate you a little bit for the value of your private information or perhaps consumers will switch from a model of

advertising-based publishing to consumer micro-payments.

So I think that's part of how to think of it. Now, in terms of, you know, you brought up these, you know, what happens under an opt-out versus opt-in, and I kind of ballparked a guess of what would be the proportion of consumers. The problem with that exercise is that it's really hard to know. That equilibrium could change very quickly. Like, right now there's a very tiny amount of people that opt out using the industry mechanism.

Now, if there were to be some huge scandal where everybody's information became available on some website, then that could change pretty quickly. So, you know, inherently it's hard to talk about those things, but I think it's important to keep those big-picture numbers like $3 or $4 per person in mind when we do this discussion.

By the way, in Europe they're considering in May 2018, as I understand it, and I'm always trying to wrestle with just what they have in mind. So I appreciate if people could clarify this for me. My impression is they want to move to an opt-in based system. And you would expect that if you ask a bunch of consumers would you like to opt in to being

followed around by 100 different advertisers, the answer is going to be go away. So that could, you know, pretty radically shift things in Europe and that would inform what's going on here for sure.

All right. Well, we'll yield the time to the next people then. Thank you.

(Applause.)

53

SPONSORSHIP DISCLOSURE AND CONSUMER DECEPTION: ASSESSING NATIVE ADVERTISING IN MOBILE SEARCH

DR. JIN: Thank you. We'll move on to the next paper by Harikesh Nair from Stanford University about Sponsorship Disclosure and Consumer Deception: Assessing Native Advertising in Mobile Search. And to make sure that our presenter would have full 25 minutes, I would request you to hold back your question unless it's just for clarification. Thank you.

DR. NAIR: Thanks, Ginger. Good morning, everyone. Thank you again to both Marketing Science and to the FTC for organizing this conference. It's really fantastic to bring these two institutions and fields together.

So this is a paper co-authored with my colleague, Navdeep Sahni, at Stanford, and this is based on a bunch of field experiments that we did with a mobile restaurant search platform. And we have two papers that came out of these experiments. One is on assessing the role of advertising as a signal, and this particular paper on native advertising gives us a sense of how deceptive native advertising is. And there's been a lot of interest in this topic. So I'll try my best to give you a sense for what we've been

54

doing.

Native advertising also has a long view throughout the century. So let me start by giving you a sense for how this issue has played out in media historically. In the 19th century, at least in the United States, most of the news media in the U.S. were owned by particular parties, and that changed very rapidly at the turn of the century as news and media oriented to a more professionally oriented journalism and journalists started emphasizing the core norms of objectivity and autonomy.

And in that business model, rather than get money from political parties, the ad supported business model was born. And in that kind of situation, in order to make sure that news was autonomous and media was autonomous, publishers instituted a separation of the church and the state that separated the business side from the news side of the media side of the business.

And the so-called separation between content that is produced by a media platform and advertising is a steadfast principle of media, okay, and has been very clearly pointed out in the previous conference that the FTC did on native advertising. That line is very fast blurring in the -- in the digital ecosystem

55

right now because of the advent of native advertising. And there are many definitions of native advertising, but one thing that we can all broadly get behind is that it's advertising that kind of matches the form, the style and the layout of the media content into which it's integrated.

So it's really content -- it's really advertising that kind of looks like content. So the line between what's content and what's advertising, it's blurring and that kind of advertising is kind of proliferating. We have a large number of estimates in the industry. So there are various kinds of very large numbers going out there. But that kind of advertising is actually the one that is gaining a lot of prominence, especially on mobile where a lot of attention is going towards within apps, in-app advertising or whatnot.

While industry adopts that format, there is a significant policy concern. And from the regulator side, the main concern is of deception, which is that consumers are harmed when the commercial nature of content is not properly disclosed.

As Ginger mentioned, the FTC has a very precise term for what is a deceptive practice. A practice is considered deceptive if it's likely to

56

mislead consumers who are acting reasonably and it would be material to the decision to buy or use the product or consume the advertisements. Okay?

So we're going to try to assess to what extent native advertising on the particular platform -- it's a case study -- is going to be materially deceptive and to give a sense for what people mean when they say something is material. It essentially means it affects their actual actions, okay, in some fashion with respect to the advertisements or to the product, which actually if you think about it imposes a high data bar because you actually need to observe actions in order to make a real statement about it.

As we all know in this room, paid search is a very large component of digital advertising and therefore assessment of deception in that marketplace is likely to be of interest and of impact to the digital advertising industry. Okay?

New regulations have come in in 2015, in the last month of 2015, where the FTC now stipulates that any disclosure in online advertising must be sufficiently prominent and unambiguous in order to change the apparent meaning of the claims and to leave an accurate impression to the exposed user as to the commercial nature of the sponsorship of the content. Okay?

If you've been following the press on this, these regulations have been controversial. The digital advertising industry has expressed some skepticism about it and a debate has been going on. Unfortunately I could not make it to the disclosure conference yesterday, but I'm sure that there are various opinions on this.

Generally industry bemoans government intervention in the creative process and believes that self-regulation and current levels of disclosure may be sufficient. In particular, the official IAB statement in response to the FTC's guidelines said that it may be overly prescriptive, especially absent any compelling evidence to justify some terms or the other. So there is really a large paucity of studies in this area. So hopefully this paper has something to say about this in one case study, and that will spur more broader studies in this area. That's the idea here.

Okay. So the goals of this particular paper are to look at does native advertising work, which is important
to establish before we proceed to see whether it's

actually very important from a deception perspective. And I guess one of the aspects of the paper is also to present a new way to assess deception without asking people whether they were deceived; instead to focus on real preference arguments alone. Okay?

Then we're going to assess that for paid search ads using a field experiment implemented on a restaurant platform. So let me do my best to get that out there and I look forward to comments and reactions.

So just to level set the audience, here are two kinds of in-app advertising from two platforms. One is Yelp, which is similar to the one that I'm going to talk about. The other one is Facebook. I will search for restaurant in the Bay area near Palo Alto, and out comes an ad for a restaurant called Bliss Pops on position one. And that's in Redwood City. That is closer to my geography of search.

And you see that Yelp reveals that it's an ad with a yellow label, which is similar to what Google used to do a few months back, and that's a nature of sponsorship disclosure, that this is a sponsored ad. On the right side is an in-stream ad within the Facebook screen for Progressive, and they reveal that this is actually sponsored by Progressive

using a gray label with the word "sponsor" as opposed to the word "ad." Okay?

So then my question, you know, is this deceptive or not? Okay. Is this is a deceptive ad? How do you assess that? Okay?

So here's a stylus picture that gets a sense for how to address that. Here is a screen shot from an app. And let's say you search for a restaurant and then three restaurants show up above the fold, and one has an ad on it. The real question is how -- how do we as researchers decide whether this is deceptive or not. Okay?

The existing way of doing this in our view has significant drawbacks. Most of the existing approaches involve exposed survey with recall. So you might be called in a random phone survey and you might be asked when you put a search last week on Google or on Yelp or whatnot, did you realize that there were paid ads shown? And if the consumer says yes, you might be asked did you realize that it was deceptive? Were you deceived? And you might say, yeah, I was deceived and we -- we have -- and the researcher may report the percentage of consumers who said they were deceived or confused or whatnot.

Now, a couple of criticisms of this approach

which are well known would be that first you're not really assessing deception in a context where the -- where the fact that they were deceived is not consequential. For example, if I'm really looking for a nice dinner with my wife and we have a babysitter at home, do I really care about the fact that this is not the right restaurant that I was recommended to?

The other one would be the exposed recall may be imperfect, and the way you ask the question may prime deception, and that's a well-known aspect.

And the final thing is that the marginal consumer to whom disclosure may change the behavior, for whom disclosure may actually be materially deceived. He or she is -- that individual is the one that we care about. There's no sense that the survey is identifying the opinion of the marginal consumer. It might be the average consumer. Okay?

So we're going to try to find a way to assess deception using a real preference argument. So we are going to construct an experimental design that is going to get at that. And that experimental design may be useful in other situations we think in which we would like to assess issues related to disclosure, more broadly defined not just in advertising. Okay?

So here's our idea. We're going to

61

randomize people into a new condition. So in the middle is the current disclosure condition, which is what we want to assess whether it's deceptive or not. Okay? We're going to randomize consumers into a condition which we call a prominent disclosure condition in which the fact that this is an ad, okay, is highlighted in a much more prominent and conspicuous way. Okay?

Now, we can think about what will be a prominent and conspicuous way, but we are going to implement a particular way of doing it, which is to highlight the ad with a border. And I'll show you exactly what we did. And we think of it as two different worlds. One is the current world, that's the middle one; and the one on the right side is a full information world, a world in which consumers fully understand at least that this is an ad.

We are also going to randomize consumers into another extreme world in which the same listing is provided of the same position here for a restaurant one, but without any disclosure that this is paid advertising. Okay? So think of this as two extreme worlds, one in which there's full information on the right and one on the left is absolutely no information on the left. So it's full deception on the left.

62

Okay?

And then we're going to track behavior under each of these conditions. Okay? Then we're going to ask whether the behavior under the full information world looks similar to the behavior under the current disclosure regime. Okay?

Now, if your choices look very different when you are fully informed versus currently, well, that means that there was deception because actions are very different. So that's very simple. And if -- so just by comparison of the current disclosure condition to a prominent disclosure condition will give us a sense for whether there's deception or not.

Okay. Now, if they are similar, stickily similar, we say that we cannot detect any evidence of deception. Okay?

Now, comparing the current disclosure condition to a no-disclosure condition, if I find that behavior is roughly the same in a world where you are not told that this is sponsored versus the current world, well, how do I -- what do I conclude from that? Well, we kind of conclude that, you know, issues of disclosure and whether or not this is an advertisement or not is not that relevant for consumers in terms of how they make actions. But if we find that -- if I

63

don't disclosure behavior changes dramatically, that also tells us that, you know, this is something that we really need to care about as a regulator because if I don't disclose behavior would be quite different. Okay?

So that's the basic idea of the design. The main advantages, it's based on real preference on actions alone. We don't need to ask consumers anything. Okay?

Question over there? Yeah?

AUDIENCE: Okay. Just a quick question. So I was curious as to do you see the behaviors changing with time? So initially I might just think I've been (indiscernible) for a long time and I know that the top one is an ad.

DR. NAIR: Correct.

AUDIENCE: So -- but, you know, over time in these two different populations, do you see the behavior change?

DR. NAIR: Yes. That's a very good question. Definitely there will be some dynamics in those and potentially some learning about the platform as a whole. We are not able to assess that, those dynamics, because for an econometric reason I'm going to assess my outcomes at a single point for the first

64

search of consumers on this platform and their response to the first search, just because of an endogeneity problem that comes up.

In the paper, we actually report what happens at the end of our experiment, which is roughly a month into the experiment, and the results that we report persist and there is some attenuation of that. But I cannot speak more than that. My basal guess is, yeah, of course there will be dynamics as people learn and understand the platform. So -- but we can't speak much to that in this paper.

So you have to decide --

DR. SMITH: Yeah, I'm sorry. So I'm curious why you're -- for the low scope for disclosure to see why you're comparing the no-disclosure to current disclosure versus no-disclosure to prominent disclosure.

DR. NAIR: You know, the way we were thinking about it is that the current disclosure to more prominent disclosure is easy for a firm to implement. And if -- in a world with full information, choices look very different. That is evidence of deception.

The one on the left might be actually more difficult for a firm to implement. Actually, I talk about how we were able to implement it because the same advertisement has been shown without any disclosure to consumers that this is actually an advertisement. So it might be hard to implement that in practice. And I was trying to tell you why such a design may be actually useful because you could get a sense for if I rate a disclosure from all the way from nothing to very high, if there are very dramatic changes, that will help us to understand to what extent do consumers care about disclosure.

Okay. So let me just skip this in the interest of time and tell you a little bit about the platform. So the experiments were implemented on a platform called Zomato, and it was implemented in 2014. Zomato turns out to be a very large restaurant search platform in many countries that were underserved by traditional search and digital platforms, in particular in India, Jakarta, Manila, Dubai, which were the cities where our experiment was implemented.

In 2015, they acquired another platform called Urban Spoon in the United States. Some of you may know about it. And so they were getting pretty big in the United States and Australia as well. But the U.S. data and the Australia data are not in our

experiment. Okay?

To give you a sense for it, Yelp is the largest local business search engine in the United States. They had roughly about 100 million to 120 million visitors in 2015. Zomato has about 80 million visitors. But Yelp is not just for your restaurants alone. They're for all local businesses. Okay?

So to understand the context of our experiment, in 2014 August when we implemented the experiment, the Zomato platform had a robust advertising market for searches on the desktop on Zomato.com. But there was no advertising on mobile. Okay?

Many thousands of advertisers would be advertising on the platform. You would put -- if a consumer puts in a search, a search ad would be shown, but there was no mobile advertising. So this experiment was implemented as part of pre-mobile test and learn methodologies for the firm. And then in August 2014 we go in and implement the mobile advertising experiments.

In the end of September, a new update was launched on Android on the Google Play store in which mobile advertising was actually included. If a consumer downloads that update, he or she is out of

the experiment and the experiment ends. So almost all the data is from in August of 2014.

Here's an example search session. This is from the Android app on which the experiment is implemented. As you open the app, you can start putting a search for a restaurant. You can use any of the pre-filled categories. For example, most of the searches -- many of the searches at least in this -- the countries that we implemented the experiment are for home delivery.

And once you bring a search, a bunch of listings show up and the figure that is in the green is the average rating of users on Zomato. And let's say Café 6 is one of the restaurants, and if you click on that listing you'll get to a restaurant page where additional information is available. So let me zoom in on that. And this additional information would involve a map of where it's located, additional reviews, you can see the menu. And, in addition, you can call the restaurant and make an order or do something else. Okay?

So it's quite information rich. And then we are going to take -- use measures of consumer activity, two measures. One is click -- whether or not you click on the restaurant, and the second

whether or not you call the restaurant. Okay?

We do not have actual orders placed to the restaurant as of this point. I don't know any phone that actually correlates in-app or online ad behavior all the way to restaurant sales. So we just don't have that.

Recently, Navdeep and I, we have audio-analyzed a large number of MP3 files where we actually listened in to about 3,000 calls that were made because we recorded a bunch of them. And we report that roughly 75 percent of these are about home delivery, making an order, because there's no real Open Table in these markets and most of it is for delivery. So we think calls is a much more important and more credible metric of actual orders compared to clicks on advertising. Okay.

The experiment was imported as an update into the app. It was launched from the Google Play app store. Any user who downloads it in one of these cities is in the experiment, okay, and then stays in it. So it's persistent user randomization over time. There's no re-randomization at the session level.

Okay. So these are the conditions into which users are randomized. The typical disclosure condition is the one in the middle where the fact that

69

it's an ad is revealed through a yellow label. Okay? The prominent disclosure condition is the one on the right where we add a yellow label to -- sorry, a yellow border to it. And then the no-disclosure condition is on the left. For instance, the advertiser Mia Bella occurs in the same position in the same location, everything remains the same, but there is no disclosure to consumers that this is actually paid advertising. Okay?

So just to clarify, there are no ads on the restaurant pages. There are only ads on these listings. So these are paid search ads. And then everything else about the listings, including the position, the nature of the content, color, everything else remains the same across these conditions. Okay?

So my full information world, what I'm calling is on the right side. The no-information world about the sponsorship status is on the left side. We also randomized consumers into a condition where there are no ads. Okay? So in this particular example, Mia Bella, the restaurant, is not advertised, but it may show up if it's relevant somewhere down in the organic listings. Okay?

There are some more details about the

70

experiment, in particular how we picked advertisers. For instance, we did not randomly pick an advertiser. We did not randomize our advertisers. We picked advertisers who actually wanted to advertise, which was important for making sure that this is data from a signaling equilibrium, which we wanted to test.

We did not want to disturb an equilibrium, okay, by picking a random advertiser who will not have advertised and showing an ad of that advertiser. But in the interest of time, I'll proceed a little bit more. And if you wanted to get more details on the experimental design, please approach me and we will be happy to talk offline.

In addition to the ad label, we also changed the way in which disclosure is included by using the word "sponsored" instead of the word "ad," okay, because there's been some questions about not just noticeability but also interpretation of the label. So we tried sponsored. But then we have a randomization into just sponsored condition, which is on the extreme left, and the condition in which the sponsor with a highlight on the second one from the left. Okay.

So there's one condition in which there's no advertisements. There are other conditions in which

71

there is advertisements with and without highlight. Yes?

AUDIENCE: I think I am seeing this right, but you didn't highlight the word "sponsored" like you did he word "ad."

DR. NAIR: Correct.

AUDIENCE: Is there a reason why you made that decision?

DR. NAIR: We did a little bit of pre-experimentation testing, and this seems to be the one that we feel that the survey said users fully understand that this was an ad. Yes. And I think there is psychologists and others who think about vision and others who have done more studies in that. And so, yeah, those are additional ways to consider it. But this is what we have done, yeah.

Okay. So there are 321 locations. A location is a five mile by five mile zone within a city. That is the unit of geography at which ads are sold on Zomato.com at the desktop. So there's -- all the randomization is at that level. And there are roughly 622 advertisers spread across these 321 locations. So it's a larger scale to the extent that we have more than one advertiser. Okay.

Okay. So, this was related to your

72

question. It so turns out that the consumers who were randomizing your condition and saw the first ad -- the ad exposure are randomized, but the set of people who came back may be different from the set of people who saw the first ad. And therefore we are going to base all our results on the responsiveness to the first ad. Okay? And there's more results in the paper.

Okay. Let me take up three of the main results. The first one is that consumers do not notice enough the sponsorship disclosures of native ads, and thus are tricked into clicking on them. And here what I'm showing you is the probability of calling, okay, relative to the typical disclosure condition. Okay? So the baseline is the typical disclosure condition, and the box represents the difference from the typical disclosure condition -- of the highlighted condition and the no-disclosure condition.

So basic punchline of the paper is that when the ad is highlighted, there is no difference. Okay? So we don't find any evidence of deception. Okay?

A second punchline of the paper is that when ads are not disclosed calls fall. Okay? We explore this a little bit more in our separate paper on signaling. We show that standard signaling models can explain that phenomenon, in particular calls to a restaurant increase in the presence of disclosure. Okay? So that's an important finding from the paper.

And there's a bunch of results in the other paper, in particular documenting that it so turns out that the better rated advertisers, restaurants, are advertising. There is -- consumers who have more uncertainty are the ones who respond more to the disclosure. And restaurants about which consumers have more uncertainty are the ones who get more bang for their buck from the disclosure that seems to be consistent with signaling.

How much time do I have? Zero, Garrett, but go for it.

DR. JOHNSON: All right. So just going back, like, how large do you think the difference would be from the highlighted exposure and, given that expectation, what was your power to detect a difference?

DR. NAIR: Yes. So there's a bunch of questions about power, okay? The -- I can tell you just off the top of my head, the power is not a big issue in this paper because the difference from the typical disclosure condition to a no-disclosure

condition -- for instance, that P value is to the order of .002. Yes? So you've got to find, like, some serious occurrences by chance in order to move that -- move that P value all the way to .05. So -- and then we have exact P values reported in the paper as well that take the power into consideration. So that's just a very quick answer off the top of my head.

I've been asked a question before, so recently we've been doing more and more on assessing power and making sure that this is not something that occurred by chance. And now P value is just too small to succumb to that. Okay.

All right. So the basic punchline here, therefore, is that we find no evidence of deception, but we find that there is a strong case to regulate because in the absence of disclosure consumer behavior looks very different. So if a typical disclosure is not provided, behavior could be very different.

We found no difference between sponsor and ad label
conditions. So the consumers do not seem to be confused by the label. And, finally, quickly in one minute, assessing the idea that if consumers click, then they continue to buy because maybe they have some

search cause or inertia or whatnot. So native advertising works by tricking consumers into clicking, and that's the way the mechanism works. I'm just telling you that we don't find evidence for that at least in our data.

Firstly, there's about an 85 percent chance of continued search after clicking on an ad. So there's lots of search happening. So it does not seem to be that you click and suddenly you buy exactly what you clicked on. There is substantive continued search after click visiting an advertiser's page.

On average, people visit about 50 to 60 listings before calling an advertiser if they call. So that seems to be an outcome of fairly thoughtful search and deliberation.

In addition, I will just read out the result. We find that much of the improved conversion for people who have been -- to whom it has been disclosed that this is an ad comes from people who actually don't click on the ad. Okay? But they get exposed to the ad, they continue searching, all within the same session, they put another search and then click on the organic listing of that ad. Okay?

So it does not seem to be that much of a lift is coming from people who click on ad, but from

exposure. So the mechanism by which advertising works in this market seems to be exposure, not clicking. Okay? And therefore we think that clicks are actually a very bad way to assess advertising.

Okay. So the punchline here is that there is very little evidence of consumer naiveté, a locking or inertia condition while clicking. And so the notion that consumers are tricked into clicking and they stick with that click does not seem to have much support in this data.

No detectable evidence of material deception, at least in this market. Choices look pretty similar to a world with full information, and ads seem to work on the basis of exposure. Some other data from Brett Gordon and Florian Zettelmeyer are doing with Facebook also seems to suggest this.

Okay. So I'll skip this. The punchline here is that in a world without disclosure we find that consumers would have gone to restaurants that were more poorly rated and had fewer ratings. So it seems that disclosure actually helps consumers.

Okay. So I just want to emphasize that we can't really speak directly to consumer welfare because we actually don't observe actual choices. But it just seems to suggest that consumer choices do not

77

change materially and the ads are more prominent. So listening to the concern for welfare losses from current disclosure standards at least in this market may be minimal. So -- and advertising seems to help consumers, okay, because of signaling.

Thank you. And I'll just put up my conclusions out there and look forward to comments.

(Applause.)

DR. JIN: Thank you, Harikesh. Our discussant is Yesim Orhun from the University of Michigan.

DR. ORHUN: All right. Thank you. Could somebody help me make this full? I'm not a Windows person. Control what? L?

Thank you. Thank you for inviting me here to discuss this paper. Let me jump in in the interest of time and really, first of all, emphasize why the design of this paper is so neat and so useful to understand material deception.

So if you look at the FTC policy statement on deception, there are three things you've got to care about. First that there was some reasonable potential for being misled. Second, consumers were kind of acting reasonably, and that at least some material changed. What does that mean? They consume

78

or choose differently because of the deception. So that is actually a choice argument. So that lends itself very well, as Harikesh explained, to a field experiment to revealed preference to ask, you know, would people have chosen differently except for deception.

Now, that may seem very straightforward to do in the field. It actually isn't because it's different than what is the first question you've got to ask. Right? So what is the counterfactual? That counterfactual, you know, you may use structural methods, but in this case actually it's not that easy even with a field experiment. The comparison should not be no ads. That doesn't make a lot of sense, right? If native ads are different than a no ad world or a different ad world, that could be because native ads are differentially effective. That doesn't necessarily mean they are deceptive.

So the question is how do we link the change of behavior? Not only demonstrate that the behavior is different, but link it to deception. And the paper does a very neat job by focusing on a very specific question. I'm going to rephrase the research question in my own words, which is do native ads mislead consumers to think that they are not ads. Okay?

79

So basically once you ask the question this way, much more precise, and honestly much more relevant for this topic, then their experimental design is really right on the money. It answers this really relevant question by putting two bookends to it. Native as a middle, the two bookends are full deception where literally you put the ad and don't tell people that it's an ad. Right? Which you don't do in experimental economics, but you can do with field experiments here. Full deception. And the other bookend is full information.

Well, for the sake of argument let's say it's full information. You may have quibbles about whether highlighting it makes it full information, but this is actionably the best the authors can do. And I actually buy it. Okay?

So those are the two theoretical bookends that they are able to implement very well in the field. In other situations, like think of Airborne's claims of -- you know, that were false of, you know, preventing you from getting the flu, you may have difficulty thinking about these two bookends. What would be a very fully deceptive advertising and what would be full information advertising isn't as clear.

But in this case it's perfectly clear and

80

actionable, executable in the field. So this is great. So they have six conditions. For the interest of time, I'm not going to go through them all. Harikesh did go through them.

The relevant one isn't the no ad condition for the reasons we talked about. If it is in effect, it's not very clear if it's because of deception.

First, I also want to simplify the design by pointing out that sponsored versus ad doesn't matter. So let's just look at this design as no ad condition, deception condition, native ad regardless if it's an ad or sponsored, and then full information condition. Okay? So basically four conditions.

And what Harikesh argued is that comparison of native ad to the full information is the way to figure out whether this ad was deceptive. I would actually also add that comparison of the native ad to the full deception condition is another way to figure out whether this was deceiving. If they're very similar, then I would say that's deceiving.

One other way of, you know, interpreting is that consumers don't care about the disclosure. But I don't think that you can pull the two apart, whether they don't care or whether they don't notice.

But in any case, these are the three main things we're going to compare. For the first sign of regressions, Harikesh didn't have the time to go into detail, so let me do that. They actually don't compare these exact three. They pool the full information, the two together, to get power. Since sponsored versus ad doesn't matter, they might have as well pooled all the native ones, which I think is a good robustness check.

What they find is actually no effect on visiting the restaurant's page. So if you just have this result, you might have thought, well, maybe they are deceptive or maybe this is not effective, but thankfully they actually have much more to say. They look at calls and they find a huge difference between the deception condition and the other two conditions.

And the other two conditions are actually insignificant from one another, not different from one another. And so they conclude that the native ad is much closer to the full information case than the deception case. That's why the bookends are so useful.

They do another thing that I think is very valuable that Harikesh didn't have time to talk about. They actually look at how the type of restaurants' consumers call changes as a function of disclosure.

So they see -- they run an interesting regression so I'm going to talk about this in detail. They look at the number of calls a restaurant gets across different conditions using restaurant fixed effect. So this is a within-restaurant -- it controls for all the heterogeneity -- the data is really rich. They can control for all kinds of heterogeneity, including search characteristics which I urge them to do, and restaurant characteristics.

So if you just look at the main effects, you might interpret this as kind of an effect of experimentation on all of the conditions on all restaurants. But, by the way, these are not advertised restaurants, these are all the restaurants.

They don't find a main effect there. It's kind of comforting because you don't actually want your experimental conditions to kind of change the calling behavior to all the restaurants, but just, you know, implemented or shopped restaurants.

But interestingly they do find that the type of restaurants the people call changes. People are much less likely to call high-rating restaurants. They're much less sensitive to rating. This was actually a little confusing to me. It's an interesting result. But why are the conditions

affecting all the rest of the restaurants which we think we are keeping fixed? It might be useful to discuss.

But in essence, what's important to take away is that the native ads is much closer to full information than deception case, even here.

Another set of results that they have which I found very interesting is to answer the question are consumers tricked into conversion. Here they're comparing deception versus disclosure. They're now lumping all four conditions of disclosure into one, which makes sense, and they're crossing it with whether the restaurant clicked on was reached organically through search and below, or by clicking on the ad.

So the paper can give more detail as to why this two-by-two answers this question. But here the findings they have on consumers are not stuck if they click on native ads. First, they show that disclosure does not impact whether somebody continues to search or not after a page visit. In general, organic arrivals search less afterwards than ad links.

This makes sense because if you went to a restaurant by an ad, you're probably -- your match value was probably not so high so you're continuing to

search. Also, ads appear at the very top and organic links appear at the bottom and the search behavior may change. It may be more likely to converge at the bottom. So there are some things going on maybe we want to control for rank.

But importantly disclosure doesn't impact. What does that mean? I actually wanted to think about these bookends again. This means that native advertising is close to deception. What do I want to make out of that? Does that mean native advertising is deceiving in this case? I don't have the other bookend. They lump the native ads and the highlighted conditions together. So one thing to potentially explore is bring that bookend back and see if it's kind of closer to full information. I was confused by this.

Another result that's really interesting is calling only increases with disclosure if the page visit was organic, not through an ad click. So their other paper is also very, very neat. They actually show that increasing in calls due to disclosure may be a signaling story, right -- is a signaling story. So my question is why doesn't the signaling story work when you click on the ads, but works in the organic links which are much lower and much less frequently visited.

So those were my kind of, you know, overview of the results. I think it's very cool. I personally took a lot away from this paper. Three things importantly that I want to re-highlight. First, the role of experimentation for identification of material deception. This idea that you can think at least theoretically of those two bookends, deception and full information, is extremely useful. Whenever it's employable, let's do it, right? This is very useful.

The elements of design in this paper are extremely clear and very well thought out. And the punchline is that the consumer response to the same ad when it's native looks similar to the full information case, but quite different from the deception case. And I want to highlight this difference between the deception condition and the native ad condition because I think that also directly speaks to deception. Thank you very much.

(Applause.)

DR. JIN: Thank you, Yesim. We still have a few minutes for questions.

AUDIENCE: This is a fascinating experiment and I think it's great. I had one question about the specific setting in which this is happening. Is this

in a setting in which ratings are easily arrivable? So we will actually have an ability to assess what the quality is because the ratings are very high. And perhaps some of the restaurants are more rated by the particular nature of the setting compared to a setting where I just don't have any information rating. So I just wanted to get your thoughts on that.

And, second, I was curious whether this is really so much more of an attention story, that when you start highlighting something you get more attention. And in a particular place where actually quality can be easily assessed, there is much less chance for deception than in other settings.

DR. NAIR: Those are two great questions. No, absolutely on the ratings, the effect that we are measuring is over and above any information content of ratings. In particular, to the extent that you believe that the advertisement serves as a signal here, in a world without ratings, the signaling value of advertising would be much higher. And in a world with no ratings -- sorry, in a world in which rates convey all information, advertising does not have any role to play as a signal.

AUDIENCE: (Off microphone).

DR. NAIR: Yes. So in terms of the effect

of ratings, yeah. So in a world where ratings provide a lot of information, the incremental value of advertising as a signal is more rated. So what we are measuring is over and above the effective ratings.

We can't say anything particular about the value of ratings in this paper because we don't randomize ratings. So where we have a conditioning on the ratings and the organic algorithm and then what we're measuring is over and above. So we randomize disclosure, but not the position on the ratings. So the paper has little to say about that.

Now, attention, absolutely I think advertising plays an important role in increasing attention. But that attention seems to be translating into clicks and exploration of the restaurants, but not necessarily into conversion. Yes.

So, for example, in the -- in the no ad condition, the listing is very much at the bottom, but in the -- what Yesim called a typical disclosure condition or the deception condition, it looks like an organic link but it's on the top. Okay?

Going from A to B is a very dramatic increase in attention because it went from somewhere down there to the top, where we find very little increase in the call rate. We do find an increase in

the click rates.

So that's -- and when moving from an ad which is provided as typical to an ad which is highlighted, also we don't increase -- we don't see a dramatic increase in call rates. So we think that attention does matter on clicks, but it's not necessarily -- just because I'm getting you into the consideration set, that does not necessarily translate into actual conversion. That seems to be the story that we see.

The closest panels I know to the literature on searches, a paper by Raluca Ursu, where she looks at Expedia but it's not for paid ads, but it's for organic ads, she sees that if you're in a higher position on Expedia, that does translate into higher clicks. But it does not translate into -- necessarily into conversion for the hotels at least in a monotonic way.

So that's why we think that clicks perhaps are not the right metric, but open to more interpretations. Thank you.

AUDIENCE: So my question was kind of related to Sridhar's question. It wasn't quite clear to me how this -- Sridhar called it attention and I call it salience. So I throw the other visual cues

89

like the box with the yellow thing would be disentangled from signaling. And my thought perhaps was it might be useful to have a setting -- and I know these field experiments are not easy to repeat, but where you have similar visual cues without the advertising message so that it would help to have the signaling story separate from a salient attention story.

DR. NAIR: Yeah. So thanks for asking that. And absolutely we do have the condition. We have a condition where the same restaurant is shown to consumers without any advertising message. And then we have a condition in which the same restaurant is shown to a consumer with an advertising message. The difference between that is what picks up signaling. And then we have another condition in which an ad is shown with and without a highlight, the difference in between that is not picking up signaling, but speaks to attention or prominence. Okay?

What we find is that if the same listing is shown without an ad, calls are lower. If the same listing is shown with an ad disclosure, calls are higher. That's why we say that seems consistent with signaling.

If the same listing is shown as an ad with a

90

highlight, calls are not necessarily higher. So that's why I responded to Sridhar's question that the additional highlighting or additional attention does not seem to be translating to calls. It does translate to clicks.

Yes, Catherine, go ahead.

DR. TUCKER: Hello, yes. So I think it's more suggestion than question.

DR. NAIR: Yes.

DR. TUCKER: But hopefully it's a doable suggestion. So as you know, Ben Edelman's got this old paper where he shows that old people, inexperienced internet people, react differently to the word sponsored and ad. And I was just thinking with your wonderful geographic data, you can actually look to see whether that's an artifact of his setting or yours and look and divide up the world into the experienced and inexperienced and see if you see any heterogeneity effects.

DR. NAIR: Sure, yes. Absolutely. So I'm aware of Ben's paper. And we will definitely look at that. But just to respond to that, the main interest of this paper is not as much the difference between the word ad versus sponsored, which is the point of his paper, but to understand conditional on the word

91

ad, if additional salience or additional highlighting changes behavior in a full information world to a small information world, and also to understand what will happen if we provide the same information but code it as an ad versus not coded as an ad.

His paper does not have a control group and talks about the difference between if an ad is coded as sponsored versus coded as advertisement, and that does change in behavior. This was not really the focus of our paper, but we are happy to report heterogeneity in that to see whether it's consistent with these results or whatnot. Yeah. And, also, his paper's results bear on clicks. We do find results on clicks. Our results on calls seem to be pretty different.

Yes, Anne?

DR. COUGHLAN: I'm kind of interested in your thoughts about the distinction between misleading and deceiving. And I'm thinking back to what Ginger said in the introduction. There's been a lot of use of the word deception that I'm not sure has been actually demonstrated here. Whether or not it changes a consumer's call behavior or indeed their purchase behavior doesn't mean that they've been harmed.

And so I don't know if anybody would like to

92

chat about that. But it would seem to be important to be precise about those words and the implications for policy and action.

DR. NAIR: Yes. I'm so glad you asked that because we grappled with this quite a bit as we were thinking through the paper. I do not think that these results translate into a statement that consumers were harmed or not harmed at all. Yes? Because harm, in my mind at least, will require measuring actual consumer welfare. And so we don't have a stance on consumer utility and we don't have a way of assessing welfare. So we don't know whether consumers were harmed or not unharmed.

But the sense in which I'm saying deception is the sense in which the FTC provides precise definition of it, deception is said to have occurred if a reasonable consumer's behavior with respect to the advertisements or with respect to the product changes. And I'm just documenting very little change in subsequent behavior when ads are highlighted and presumably people understand that they are actually ads.

Now, does that -- that does not translate into a statement about harm. I don't know because I'm not in a position to measure welfare. We tried to show that it does not seem to change in harm because with disclosure compared to a world with no disclosure, people seem to be going to better restaurants which have higher ratings, and to restaurants with fewer ratings. So it seems consistent with signaling. It may not translate to harm, but without taking a stance on utility or measuring welfare, per se, I do not know. I would ask the FTC folks to tell us how we -- I should think about it.

DR. PAPPALARDO: I have a related question, which is if you don't test the effect of disclosure on consumer comprehension as part of the experiment, then how do you know that the consumer was misled to their detriment?

DR. NAIR: Correct. So we don't really have a way of getting inside consumers' minds to the extent that we would like to. And we think that ways of asking people subject to the usual Heisenberg critique, the asking changes when they complement and how they report to that.

So I -- all I can offer you is what patterns of consumer actions. And the action seems to be of deliberate search, not of knee-jerk reaction, of substantial consumption of listings prior to calling,

and of a responsiveness of reactions to ads that seem consistent with the theory. And most of these actions don't look like knee-jerk reactions, and they seem consistent with people really understanding that what they are seeing is an ad.

And in a setting where ads are made more salient, they don't seem to be behaving very differently as well. But in a setting where ads are not at all shown, they seem to be working very differently. All of this seems to suggest that people are comprehending. But I don't really ask people whether they comprehend it. In fact, we are critiquing that style of assessing the advertisements.

Yes, go ahead.

AUDIENCE: Hi. This is from the legal perspective.

DR. NAIR: Yes.

AUDIENCE: But your statement about, you know, what's deceptive under FTC law and if behavior doesn't change then the ad isn't deceptive. And I would take issue with that because the issue on native advertising is how much weight or credibility consumers will give the message. And that's why the FTC has said it should be identified as advertising.

The consumer can understand that it's an ad

and still want to engage with the ad.

DR. NAIR: Right.

AUDIENCE: But it may change, you know, how they perceive what's being said. So it's not that if consumer doesn't click that means they weren't -- that they were deceived or were not deceived.

DR. NAIR: Correct.

AUDIENCE: It's the weight of credibility of the message, not just pure engagement.

DR. NAIR: I am in agreement with you. And I think the -- what our paper is trying to document is that the translation of that change in credibility in response to disclosure is in a positive way to the restaurant. In a world with disclosure, consumers are actually going to the restaurant at a higher rate. They're calling the restaurant at a higher rate. Okay? And that's all we're trying to say, that when consumers see the same listing framed as an ad, the credibility of the restaurant increases.

Yes? And that seems to be suggested with signaling, and this is real data that just documented that.

AUDIENCE: I just want to take on your Heisenberg example. We don't stop physics from doing measurements, nor should it stop social sciences. I

want to give you a theory of how reactivity occurs.

DR. NAIR: Yes.

AUDIENCE: And particularly in the case where many of your concepts like attention are actually very easily measured not in a field experiment, although increasingly with very cheap iTracking, $99 in a pin -- you know, a little pin-shaped container, or lots of other techniques that can be. So I want to push back on this statement that you're trying to go against that kind of measurement.

DR. NAIR: No, not necessarily as a substitute. I didn't mean to say that this field experimental agenda is a substitute for that kind of measurement. But I wanted to say that it's a complement to that kind of measurement. And I do believe that the actions of consumers when they are actually engaged in the record search for a goal that is very important. Let's say dinner with the family on a Friday evening, when it's really important to find the right restaurant, they could defer a little bit from the in-lab situation where the actions are less consequential.

I think if we can find a way to measure consumer beliefs and consumer information sets and attention in the field in a way that does not disrupt things, that will be incredibly valuable and we should potentially find a way to combine that. Right after this talk I'm going to come to you and ask how should we do that in the field, and that will be great.

AUDIENCE: Okay. So I think this is the last one. Okay. So this is actually a followup to Anne's comment.

DR. NAIR: Yeah.

AUDIENCE: So I think maybe the conditions she -- we were just talking about it.

MD. NAIR: That's fine, yes.

AUDIENCE: So I think the condition that would be nice to have -- imagine having 20 different listings and one has an ad next to it. So one way to think about it is this is a form of disclosure. Another way to think about it, this is a form of salience. It's calling -- you know, it's kind of calling attention to that ad -- to that listing. And it is also, you know, an ad disclosure, but it's also -- you know, sort of gets your eyeball to go there.

And so it would be nice to have another condition that had something like that, a type of salience, but it didn't say ad. Like, for example, it had a star.

DR. NAIR: Mm-hmm.

AUDIENCE: Or just had a box without the ad symbol.

DR. NAIR: Right, right, right.

AUDIENCE: And so I was wondering if you had that -- okay.

DR. NAIR: I see that data and the short answer is no. We'd love to have it, but, no, we don't have it yet. Thank you.

AUDIENCE: And so just kind of -- so do you -- so do you feel like in light of that, so another way to interpret the salience is to say, you know, this shows that salience -- you know, that salience is a good thing, as Sridhar was saying as well. So would you make a strong statement about disclosure?

DR. NAIR: That's right, yeah. I don't interpret it that way, but if you would like to interpret it that way that would be fine. But we don't think that salience is what's driving at our attention. Our consideration set is what's driving it. But rather than belabor the point, let's chat offline.

(Applause.)

DR. JIN: Thank you for the engagement. It seems like we have underestimated your willingness to discuss. We're a little behind schedule. We'll take

a break right now until 10:45. Thank you.

(Brief recess.)

SESSION TWO:

THE BENEFIT OF COLLECTIVE REPUTATION

DR. JIN: Hello. We'll start the second session on papers. Aniko Oery from Yale University will talk about The Benefit of Collective Reputation.

DR. OERY: Thank you. Yeah, thank you so much to the organizers for putting together such an awesome program, and also for giving me the opportunity to present here. I'm very excited that we have a session with theory work. So I'm an economic -- I shouldn't say economic theorist. I'm a marketing modeler theorist. I don't know how we call us. And so we -- I will have less to say in terms of quantitative results, but hopefully I can give some qualitative insights that are relevant for regulation as well.

And I have to apologize actually also to Anthony because we changed or added a bunch of results that are more relevant maybe for regulators. And so I wanted to -- I will focus a little bit more on that in this presentation, and I also apologize, but maybe that's good for those people who have seen this talk already, which I think there are some of you who have seen this talk.

So this is joint work with Zvika Neeman, who

101

is an economist at Tel Aviv University, and Junju Yu, who is an amazing student at Yale. And -- well, okay. And let me now jump into what we think about when we talk about collective reputation.

So there are a bunch of different types of questions that we can -- that collective reputations can help us answer. One is agricultural appellation. So here, for example, if you think about a brie cheese or a Bordeaux wine, if you go into the wine store you might not know exactly which vineyard the wine comes from but you know Bordeaux, you have some idea about the quality of a Bordeaux wine. And similarly if you buy brie, you know, you have an idea about the quality of a brie, but you might not know the exact brand of the cheese.

And so there you have some -- those cheese companies basically collectively build up their reputation or have a collective brand in their agricultural appellation.

And maybe more importantly for regulators is the country of origin application. So, for example, TAG Heuer or many other high-end Swiss watchmakers really put very prominently on their ads Swiss-made, or German manufacturers of cars put on. So if you have a Volkswagen, power of German engineering. Here

102

-- I put this example because I think it's a nice example where you can see that even if you have a brand, if you shirk and don't keep investing then bad things might happen and the reputation might suffer from it.

And this is a very important feature of the model that we have that we really think about reputation as something that you can manage, and that might also deter you from -- but there is a moral hazard problem that might lead you to no investment or to shirking.

Another -- but then on the other hand in some other industries, we observe firms not really emphasizing it so much, and also it depends on the country that the company is from. So, for example, Bosch doesn't really emphasize the "made in Germany" so much in their ads. And on the other hand, Chinese manufacturers, there are, like, webpages where you can find Chinese manufacturers of some parts advertising together.

But then also the question is do they really want to emphasize made in China, let's say, or should the regulators say you have to emphasize made in China. And so that's kind of a question that I added to the paper after -- yeah, in the last couple of

103

weeks. So this is a pretty new paper and still a work in progress, even though now I think we have all the results together at least.

And so -- so the country-of-origin labeling is a big issue. I think Ginger also mentioned it in her presentation at the beginning. And there -- it's not clear in which industry we should regulate it, what are the consequences of it, is there actually -- does it help the consumers or does it maybe even hurt them?

Okay. So the way we tried to model it, or our contribution is that we think of a country of origin as a collective brand, and we think of a collective brand as something that creates value for a firm and therefore enforces a firm to invest into the production process of the product.

And then we also distinguished between two types -- two very different types of industry. So one is industries where we care more about quality control, and one is where we care more about investing into an exclusive technology.

I'm always not sure -- so I can also step back a little bit and the mic will still capture it, or -- okay, good. Because I'm just standing here.

So the research questions that we tried to

104

answer in this paper is on the one hand what are the fundamental differences in reputation building or brand building if we do it by ourselves, if we have an individual brand, and if we have kind of a collective brand like country-of-origin or appellations, for example.

But then I think what is more relevant maybe for the audience today is in which industries and countries is country-of-origin legislation or labeling -- sorry, country of origin labeling socially optimal, and when is it actually -- when does a firm actually want to label the country of origin and when does it not want to label the country of origin? And the gap between the two will basically capture who wants to regulate this. So we want to regulate it if it's socially optimal, but not optimal for the firm.

And from a theoretical perspective, so the way we think about it is that there's a classic model by Mailath and Samuelson about reputation building. And the difference between individual reputation and collective reputation is the following: So on individual reputation, each firm sells under its own brand name. So the customers know exactly which product has been produced by which firm.

But on the other hand, for this brand we have fewer observations so there's less output produced by that brand. And if we have a collective reputation on the other hand, now if you buy a brie you might not be aware -- you have an idea about the brie, but you don't know exactly whether your idea about the brie is generated by this particular producer of the brie, or whether it's generated by some other brie manufacturers.

So this is a weaker signal that you can get about the brand value. And, on the other hand, you have many, many more signals. So as a manufacturer or as a firm, there is some free riding going on. So you might -- that might also -- you might think, okay, why are collective brands beneficial at all then for incentives? Why do we want to have collective brands? Because signals plus free-riding problems seem to be something that -- from a welfare perspective.

However, what we would like to focus on in this work is an idea that was first introduced by Mailath and Samuelson that there is a moral hazard problem in the context of brand reputation. And this model as a problem comes from the fact that investments are not observed by the regulators or by consumers. So you don't observe the actual investment, but you observe the quality outcome of the

product. So you observe the Volkswagen car, you see that it breaks down or that it's not as energy-efficient as they claim it to be.

So -- but as I said, reputation is an asset. And the nice thing about the model that they introduced is that you can really manage reputation -- but what this leads to is that if reputation is, for example, very high, you might want to milk your reputation and shirk. So that might be the case for Volkswagen that they were just so overconfident because their reputation was so high and they just thought they could get away with shirking.

And on the other hand, if your reputation is very low, you might just give up. So the question that we ask is when does an equilibrium exist in a game-theoretic sense where a firm really wants to invest in every period.

And I just want to give you a toy model to give you the main idea behind this and where this tension comes from. But I won't go too much into detail of the theory behind it because the model is quite -- yeah, it's a -- there's a lot of details that I will have to skip today in the interest of time.

So imagine there's a firm that lives for a certain number of periods and a firm can be either

competent, so it has the ability to invest into the technology or into monitoring with some probability, or it's incompetent otherwise. A competent firm can invest and increase the probability of producing a good product from a low probability Pi L to a higher probability Pi H, whereas an incompetent firm who just doesn't have visibility to invest always produces a good quality product only with probability Pi L. And importantly the investments are not observed by the market.

And in every period you have some customers arriving and they see the history of realizations of the product and then build some beliefs about how good they think the brand is actually. And then based on that, their willingness to pay will be determined. And if the quality -- for simplification we normalize everything to the value of a customer, of a good quality product being one and of a bad quality product being zero.

So then we get a very simple equation for when is it optimal to actually invest. It's optimal to invest if the increase in the probability of producing a good quality product is bigger than the cost of investment. So this would be just Pi H minus Pi L is greater than C.

And in this type of model, in the last period if there is no future and you don't -- you know, you die tomorrow, a competent firm would never want to invest because there's no value of reputation at all.

But the fact that you don't invest in the last period means the customer -- it's not useful for the customer at all to know that somebody is competent or has the ability to invest because they know that even if you have the ability, you will never do it because of -- because you don't have incentive to do so. And hence the whole thing will unravel and even in previous periods you don't want to invest at all into reputation, or into your brand.

And so this is like the classic moral hazard problem, but a little bit more dynamic and dynamic setup, and this will lead to no investment whatsoever and everybody just shirking all the time. So that's just a very extreme case that we are thinking about.

Now, of course you can say, okay, if you have long life firms and there's no final periods, then the whole thing might be alleviated but still intuitively there is something called the discouragement fact which you still get in a dynamic setup that once you have very high reputation, you don't -- you want to milk your reputation and you want

109

to milk your brand value, and if you have a very low reputation or your brand value is very low, you want to just give up and stop.

And the question is when can -- when might there be some potential value of collective reputation or a country of origin labeling in order to alleviate this problem? Because it will bundle together signals or you cannot distinguish as well between different producers and this might actually give people some -- firms some commitment value to keep investing.

So the main point of the paper or the main idea of the paper is that country of origin labeling might help against a moral hazard problem. And then another nice feature is that we can really say in which kinds of industries it would help.

So one type of industry is if we have exclusive technologies, which means it's very hard to actually produce a good quality product, so it's really about innovation. So if you are an incompetent firm you cannot produce amazing cars. But if you're a good type, then you can -- you have the ability to produce a good car if you invest. So that will be the case where an incompetent sub Pi L is equal to zero. But in these industries, collective

110

reputation can really be useful only if you are at a very high baseline reputation. So you can think of it as very developed countries where you might want to -- yeah, where the commitment value is high.

On the other hand, if we have more quality control issues where everybody can produce good quality products, but if you shirk, you make mistakes and this will lead to the product not functioning very well. In that case, the commitment value of collective reputation or collective brand or country of origin is high in countries where the baseline reputation is relatively high. So if you think of maybe some developing countries where you might actually -- yeah, where also maybe a regulator might want to enforce a country of origin label.

And then there's another thing because now so far I've only talked about the social benefit of collective reputation or country-of-origin labeling. But there's also -- of course, now does a firm actually want to advertise it by themselves? Because if they wanted to advertise it by themselves, then there's no point in regulating it at all.

And so here we have on top of the moral hazard problem, we now have an adverse selection problem in the sense that if I think -- so that's the

111

classic lemons problem. I know that there are some good firms, some competent firms, some incompetent firms. I don't know -- there's a certain fraction. And so the willingness to pay = for a product might be lower than the cost of investment, even though the benefit of the investment is higher. Because part of the surplus just goes to the bad firms because I cannot distinguish as a market between good cars and bad cars, good firms and bad firms.

And so this will create basically the gap between the socially optimal decision of -- or the socially optimal investment decisions and the investment decisions that maybe firms might actually make if they choose to brand together or not.

Okay. So how long do I have? Until 15 past? Ten more minutes. Okay.

All right. So I will now go a little bit through the model because I would like to give you a flavor of what is going on here, and then I will talk a little bit more about the intuition.

So, again, we have an infinite rise model, we have a long-lived firm, the incentive is really that the competent firm can increase the probability of producing a good product at a cost C. And here

112

what's the main -- so in every period a new buyer arrives, it produces -- it sees a good quality product with the probability that is determined by the investment in the past period, and the firm just makes a take-it-or-leave-it offer to the customer, which means it extracts basically the entire surplus from the customer. So the customer just pays whatever he believes the product is worth. And this is what creates these reputational incentives and creates this reputational concerns and the whole dynamic problem.

So it's a very rich dynamic programming problem. We have the -- yeah, we have discounting and the value of the firm is the price minus the cost. So it's a very standard problem. And the long-term tradeoff that you can see here already is that your have to invest today, but the benefit from your investment is only captured later in the future. So the benefit of reputation is a very long benefit, and people can only collect it later, and that's why we get this tension between the socially optimal investment level and the investment level that can actually be achieved in equilibrium.

And the question -- and now formally mathematically what we do is really to think about

113

when does this reputational equilibrium exist. And there are some difficulties in the sense that we need to make some modeling assumptions in order to reach that. And maybe I'm actually going to -- I'm going to skip part of this and go to the intuition here.

So one just followup thing, and I think this also makes into an intuitive sense, is that a reputation equilibrium can only exist if the cost of investment is relatively low. So if the cost of investment is a little bit too -- is a bit too high -- it will be socially optimal to invest, then a reputation equilibrium might not exist despite it being optimal to invest from a social perspective.

And, again, I want to now talk about these two extreme cases. So when is it -- sorry. And the comparison between collective reputation and individual reputation stems from the fact that this cost level at which you can guarantee the existence of these good equilibria might be different, and in the collective case it might be higher than in the individual case. Okay?

So now we have these two very different setups. So the reason why I said, okay, sometimes in

114

developed countries or countries that have very high reputation to start with, why can't exclusive technologies, collective reputation help if the following: So -- because individual reputation, individual brands feel very strongly.

So if you have high prior, so we think firms are very good with very high probability because it's a country that has a very good reputation, baseline reputation, then after seeing a good signal or seeing, oh, Volkswagen produced a really good car, you believe that this firm is actually a good firm, it becomes extremely high because you know it's super hard to produce a good quality car, and an incompetent firm would never be able to do so.

And this basically -- so the reputation, the brand value, becomes extremely high and the firm's incentive to invest deteriorates because you just want to rest on your laurels.

On the other hand, if we had a collective reputation or a collective brand, then this whole effect would be alleviated by a lot because now even if you produce something good, you're not sure whether it was produced by Volkswagen or by Mercedes, so maybe it's actually not Volkswagen that has this amazing technology. But you don't remember which company it

115

was from that country.

And similarly you can make the argument if you have very low ad reputation and quality control, you learn a lot if somebody fails. Because everybody can produce a good product, but if somebody fails then you know for sure that somebody did not invest or is just incompetent. And, again, there you get super discouraged once you observe a bad outcome, and hence, again, your incentives to invest are deteriorated.

So that's why -- that's the connection or the intuitive connection between the two. And now if we have a collective brand or country of origin, then you cannot -- the signals are not as strong so you cannot detect it. So you're less likely to reach those extreme beliefs and have incentives to milk reputation or to just give up.

So this is just a summary of the results that we have. So depending on the baseline reputation, which would be kind of the reputation of the country of origin, and the different industry types, we give different predictions. So, for example, you should -- Swiss watches have a very strong incentive to brand together while -- or to emphasize the country of origin because Switzerland

116

has a maybe very high baseline reputation, whereas maybe some manufacturers of parts in Switzerland might not want to emphasize the country of origin so much.

Okay. And what are the incentives of firms now to invest? So that comes back to the lemons problem that I was talking about before, the adverse selection problem. So now formally speaking, an adverse selection problem is really that your willingness to pay might be very low because of the probability that a firm is a good type, it's very low, you just want to -- your expected value is so low that you don't want to pay for it. So you don't want to -- it doesn't make up for the cost for the firm.

So basically this commitment value of country of origin is not internalized by the firm's themselves. And so that's basically the case where there might be some value of regulation and where maybe the Government or the legislator might want to force firms to label the country of origin. Of course, there are many other reasons why you want to label and I skipped a little bit through that slide.

So I think most of the regulation is in the food industry or where you might protect the customer for different reasons. But even there it is about

117

quality control, and I think it's important to think about the incentives of the firms to actually keep on investing and not being discouraged for these reputational reasons.

Okay. So the takeaway for today is really that, first of all, collective brands and individual brands work very differently. It's not straightforward to think about how to model these two, and we tried to use the very classic setup by Mailath and Samuelson in order to so. We can distinguish between two types of industries that have more exclusive technology versus that -- where quality control is more of an issue and can address in which types of countries you might want to regulate one or the other. And we can also maybe explain a little bit why we observe so much, emphasis of country of origin for some products versus for others.

Well, and then there's also an adverse selection problem on top of that. So if the baseline reputation is very low, this adverse selection problem becomes particular high. So in particular for maybe more developing countries, regulation might be useful.

But, of course then also regulation can be harmful if it is not socially optimal actually to

118

invest. So you have to be careful there as well. But this gives you a framework to think about it a little bit.

All right. I think I'm -- oh, I have still one minute, but I think I will stop here. But if you have questions now already, otherwise I would let Anthony take over.

(Applause.)

DR. JIN: Thank you, Aniko. On a related note, FTC does play an active role in the regulation of "Made in the USA." So I will turn the floor to Anthony Duke from the University of Southern California for discussion.

DR. DUKE: Okay. Thank you. It's my pleasure to discuss this paper. I'm going to focus my comments -- I'm going to tell you a little bit about how I interpret the paper, what I got from the paper, talk about its contribution and how I see it contributing to the broader base of knowledge on reputation. And then I'll offer some critical remarks at the end, maybe some things to think about for future research.

So the paper focuses on a common phenomenon. When we talk about country or region of origin or agricultural appellation, I was thinking perhaps even

119

to some extent where collective reputations occur in franchises. I don't know if the model applies directly towards issues of franchises, but there is a collective reputation issue going there.

And then Aniko has a nice example in the paper where she talks about two drivers who belong to the same platform. I was thinking in this sharing economy if we have individual entrepreneurs or individual businesses under a platform, the extent to which they have a collective reputation and maybe milk off of the overall brand.

What makes this so interesting is we typically think of country of origin or region of origins with a collective reputation as being high quality. All right? And so this immediately suggests, well, then there's an incentive to free-ride on this collective reputation. Right? This is the lemons issue that she's been talking about.

And so this creates this tension in the whole paper is that there's a high quality and there's an incentive to shirk on quality. And so how do we resolve this tension? When is there going to be investment by these firms?

And so the question really is about how does collective reputation form and when does it lead to

120

higher quality. So she's looking at the investment decisions of these firms in a collective -- in a collective group or in an individual group and how they differ. So I see the research objective. The model -- and I'm going to highlight the key features. And when I say the key features, these are really what define the model. And it's done in a very thoughtful way. And so, in essence, I think these features really are well chosen.

There's dynamics, of course, because you invest now to free ride later possibly. It's a five-period model and the basic model is a five-period model. And I think that's a nice way to look at it. There's a sufficient history, two periods in the past, two periods in the future; there's a short run and then there's a long run. And five periods gives you that, and I think that's a nice feature.

There's random consumer match so there's no competition. That's by design. We want to keep competition issues out of here so we can really focus on the belief formation and the reputation formation.

And there's incompetent firms and competent firms. And only competent firms can achieve high levels -- high outcomes, good outcomes. And the fact that there's income in firms here means that consumers

121

cannot perfectly anticipate quality. And that's what you need to really study the reputation issue because reputation lies in the consumer's mind. Right? That's really what we're after. And so this is a nice way to do that. And there's some other features. There's no monitoring, for example. We might think of monitoring as a way to deal with this. But we want to figure out how reputations can form without those sort of techniques, and I think that's a nice aspect.

What is the basic results of this model? Well, first of all, let me describe how they get the results. They look at -- they focus on one type of equilibrium of reputational equilibrium, and they're looking at conditions, minimal conditions in which you can support equilibria in which everybody invests all the time. Okay?

And then they compare individual versus collectives and what are the minimal conditions in terms of investment costs that sustain this equilibrium. And then they can compare these two conditions and say, okay, when is collective reputation perhaps more likely or is there a larger scope for that type of collective versus individual.

And the basic results are as given in this table. So you can think of this in two dimensions,

122

all right? The base reputation, high or low. So this would be perhaps whether it's, you know, French wine or something that would be a high base, and then she compares these two cases between exclusive knowledge and I think exclusive technology as well where it's easy to detect failure. And then there's quality control and it's easy to identify good behavior. And I've put some -- just pulled off-the-cuff examples of where these might apply.

But to get a sense of what this meeting -- what this -- these results say may be more holistically -- and I hope, Aniko, you don't cringe at this, maybe this is too simplistic, but think of it this way: So in these two dimensions we can talk about initial reputation on the left, on the vertical dimension, and on this horizontal dimension whether cheating is easy to detect or being good is easy to detect. Okay?

And so the collective reputation occurs in these two corners, and they exploit either a shadow of a doubt or what I like to call a shadow of a doubt with a benefit of the doubt. And so let me elaborate a little bit on that.

So the shadow of a doubt is if cheating is easy to detect, the benefit to the collective is that

123

there might be -- there's some probability that there might be an incompetent firm among us and consumer beliefs from this way -- they say, okay, well, you know, I see good behavior in the past, well maybe I see bad behavior or maybe I see a bad outcome, but I have a -- but there's -- let me say it this way. There -- if cheating is easy to detect, right, and I'm a competent firm, if I shirk it's going to be easy to detect. And then beliefs, consumer beliefs, will react strongly to that because they might expect that there is a competent firm in there cheating. Right?

And so this is like this -- this is supposed to be the carrot and the stick. This is what keeps the competent firm investing, because I know if I don't they might think I'm incompetent. And that's the shadow of the doubt that keeps me in good behavior.

On the other corner is the benefit of the doubt. So if this a low initial reputation, if being good is easy to detect then consumer beliefs are very sensitive to good behavior because they don't expect much. There's a high likelihood that these firms are incompetent. So the benefit of investing and getting a good outcome is very high because I can change consumer beliefs. And it's the benefit of the doubt

124

because the consumer knows that, well, there could be incompetent firms there; oh, but he's shining out there because he invested. Okay. And so that's how I interpret your propositions one and two.

So they have some additional results in the paper that she didn't get a chance to talk today in the short presentation. She talks about arbitrarily long memory. And this is where history of good outcomes may be observable, or bad outcomes, and what this tends to be good for collectives.

And what I like about this result -- and I know this is a new version and you'd put that into the appendix, but what I actually like about the result, it might help to explain the strength of some of these older CEOS, I mean, in Europe where they go to great measures to talk about or to protect, you know, regional names from being used in other contexts. Like scotch can only come from Scotland and champagne can only come from Champagne, France.

And then there's some results on brand formation about when do you want to join a collective brand. And the cool result about this is sometimes you want to include a competent firm. And this is also to keep the benefit of the doubt or to keep the shadow of the doubt active and helping and encourage firms to invest when they belong to a collective.

Some critical comments, what does this contribute? Well, there's a good bit of literature on collective branding, co-branding, umbrella branding, guild branding is a name I just came up with. But, you know, what these papers typically do is they focus on a situation where reputation is already established and then what do you do with that.

This paper and the point of departure I see is where these reputations come from. I know there's some work in micro-theory, but not from a collective standpoint. And so this really gets into the microfoundations of where beliefs come from for a collective reputation. I think this is a nice contribution.

And obviously it has relevance for marketing should firms join and how to regulate, which Aniko talked about today.

So the positives, it's a meaningful research, it's carefully constructed and it provides some novel insights. Going forward, I broke up my going forward into T plus one and T plus two, just like in the paper, with T plus one meaning maybe think about for the -- as this paper develops and the T plus two more forward looking.

The papers are a bit tedious to read, but it's worth it because of these insights. I have some thoughts on maybe how to get around that. Maybe it's not possible, but we can discuss them later.

The brand formation, which I think is a good direction to go and is a nice start, so when do firms -- and the brand formation is when do firms join a collective and when they don't. My concern a little bit, or at least I think the one concern you'll need to think about is whether the decision to join is potentially informative for the reputation. Okay? And does that decision -- you've probably already thought about that, but when I was reading I was thinking that might be something you might have to deal with.

I think you're safer basically on the regulations side, on the labeling versus not labeling, and whether regulatory bodies want to grant that sort of collective demarcation.

Going forward, two plus two if you will, I think this brings up new questions about when you have a collective, what are some of the incentives to invest when outside firms try to sponge off of a collective reputation. Like, I don't know if you heard the story recently about this guy who's

importing grapes from Napa Valley to Texas to make grapes, or you can send your recipe to Belgium and a monastery there and they'll make the beer for you and ship it back and then you can say it comes from Belgium -- from a monastery in Belgium and sell it in the U.S.

But I think I'm out of time so I'll stop there. And I just want to say it's a nice paper with lots of cool insights, and I look forward to seeing the next version.

(Applause.)

DR. JIN: Thank you, Anthony. We can take a few questions.

AUDIENCE: Could you tell us a little bit more about how free-riding works in the model? Does it affect either of these cases more than the other?

DR. OERY: Yes. So it definitely helps the individual case. The individual brand benefits because it's a problem of having too many firms there, and so you kind of want to free ride on other people's investments as well. So it kind of -- yeah, it goes -- and I think we focus mostly on the case where collective reputation can be useful because when do we want to maybe enforce labeling of country of origin. And that's why I included it, because even

then it has a great value. We wanted to make sure that there we were robust to this.

AUDIENCE: So there's some -- I think some of the reasoning that firms are interested in country of origin and appellations of origin is for competitive reasons as well as reputational reasons. Have you -- obviously for trackability you assume no competition here. Do you have any more thoughts, though, on how that would affect your results?

DR. OERY: It would -- no, I don't know at which direction it would go. We have thought about it a little bit and it just becomes a mess once you -- because then you have to make assumptions about, okay, how does reputation really enter the firm's profits, because then, yeah, you have also these competitive concerns so the pricing becomes much more messy. Right now we just assume the firm can extract everything from the consumer.

So the consumer in our model doesn't get any surplus, basically. And we really want to purely focus on the incentives of the firms. But it would be nice if we can find a nice way to model it, that would be great. I don't want to make statements about how it would affect it. But, again, because we have so many differences, we have collective versus individual, and then also difference between industries.

DR. JIN: Any more questions?

(No response.

DR. JIN: Okay. Thank you.

DR. OERY: Thank you so much.

TAILORED CHEAP TALK

DR. JIN: Our next paper will be presented by Pedro Gardete from Stanford University about Tailored Cheap Talk.

DR. GARDETE: All right. Thank you very much for having me. It's a real pleasure to be here. This is a co-authored paper with Yakov Bart, who is a professor of marketing at Northeastern University. And he's teaching marketing as we speak, so he couldn't be here. He's very sad about that. But I'll do the best I can without him.

So the title of the paper is Tailored Cheap Talk, and the starting point for the paper is the fact that lots of matching markets rely on communication to make those matches occur. And a relatively new trend that's happening is this process of tailoring. So the fact that I can acquire information about consumers or whoever it is that I'm trying to persuade for a given behavior and use that information to customize my communication to those consumers.

So this paper is going to investigate the role of communication and matching. I'm going to talk a little bit about the process of data collection and whether I should want to disclose that I'm collecting this data, for example, to consumers. And then we're

also going to talk a little bit about privacy and welfare implications.

So we're going to consider a very simple model. So competition, again, between two parties. There is a persuader and there's a receiver, and the persuader has the ability of sending a message to the receiver. Very, very simple.

Moreover, the persuader also can collect information about the receiver's preferences before sending a message. There's going to be two extra assumptions. So the receiver, at least to start with, the receiver is going to be able to observe the quality of the information collected by the sender. And I'll qualify this a little later.

Morever, the receiver is going to be strategic. So the receiver understands that whenever she gets a message -- I'll be using a male sender, female receiver, just to make it simple. So whenever she gets a message, she may think, well, that's great, this is great for me. On the other hand, it may be too good to be true. So we're going to allow that. The receiver is going to be strategic. I haven't introduced my receivers into this, but whatever happens to strategic receivers it will probably work a little worse for naive receivers.

So we feel our model applies to a number of matching markets, whenever you have one of the parties trying to induce an action from the other side of the market. So, for example, in the job market you can think -- in the job market world you can think that a job applicant wants to persuade the potential employer to hire him or her. And so in that case the persuader is actually the applicant to this market.

And this persuader also has the ability to acquire information on this company. And, moreover, there's a job post so that is also relevant information. And so there's information acquisition from the applicant's side. On the other hand, there's information disclosure on the hiring side.

The bidding market, if you think there's a persuader and the persuadee, then one of the parties is trying to convince the other of very high match values. And, of course, you know, if you were thinking of online dating, of course we'll use the profile of the other person as sort of information they can use.

And the first person actually can devise when they're designing their profile, they're also understanding that this information can be used for persuasion. School admissions is another case, that is very similar to the case I just talked about.

And then if you think about relations between companies and other companies or relations between companies and consumers, and in procurement contracts, sales, advertising, we have situations where a company is trying to basically persuade a potential client that it has the right product or the right service to satisfy their needs.

Given our setting, I'm going to add a couple of comments about advertising in particular, although the model applies to other settings as well. One way to think of this model is maybe not what's happening right now today in advertising, but in a sense we're peeking a little bit into the future and looking at the consequences of some trends that are occurring right now.

So if you think of what's happening in terms of information acquisition and how easy it is to get information about consumers, that's just becoming easier. I have here a number of points of realtime acquisition of consumer data, you can get this data across multiple channels, across multiple devices. It's never been easier to store and acquire this information than store it. And there is this whole emergence of a new industry. This industry has

existed for a while, but now has expanded tremendously with data brokers.

So the first trend is it's easier and cheaper to collect better information about consumers. And the second trend is the ad delivery technology. So I never know exactly how many milliseconds advertisers have to bid for a particular impression. But in that set amount of time they can also decide the ad copy that they would like to deliver to a different consumer.

And so this is also a highly automated process right now and we're in the -- I think right now in the situation where the technology is -- the trend is to connect these two. So more and more we're seeing the information acquisition about a particular consumer being used to -- in terms of -- in the form of a dynamic creative to be used to give this consumer a different message.

All right. I also want to situate the paper a little bit in the literature. And so we will be on this top row in the literature. So we'll be looking at persuasion via cheap talk, which means that we're going to allow the sender to engage in misrepresentation. This could be lawful or unlawful.

So you can think of persuasive puffery.

That could be one form. Or you could think of the sender trying to use the infinite possibilities of language to imply certain things that, you know, are not strictly said but they're meant. Or you can also think if you talk to Upender and I and our purchasing -- car purchasing decisions, you can also just think of situations where with 99 percent probability the salesperson told you something that may not exactly have been true. And Upender was immune to that. I just fell for it. But that's -- that can happen.

So you can think of this spectrum. And this -- it's good -- in this model, the receivers are still strategic. So no one is being fooled, but despite that there may be issues about communication and persuasion.

So you can think of this paper as uniting this literature with the one on the top right corner on information acquisition and one-to-one advertising. So we're basically giving a particular mechanism of persuasion to this literature.

I also want to contrast it with two other literatures that occur. One is on the bottom left corner, it says Persuasion through Disclosure. So in that case it's a little different because I can decide whether to disclose or not certain attributes, but if

I disclose an attribute I have to be 100 percent correct about that disclosure. So we're not going to be looking at those cases. And those cases have received more attention in the past.

Another case is the case of Deceptive Advertising. So it's this very old intuition that I may just use the costs of the message or how much my investment was in advertising to say that, well, I must have a great product otherwise I would not advertise as much. And we're actually right now working on that. So I'll have a little bit to say. We're actually incorporating the cost of advertising into this and we can replicate some of the results but have also some intermediate results as well.

So I won't have enough time to go through the specifics of the model, but I wanted to give you an overview of what's going on. There's two parties in this model, there's a sender and the receiver. And they're going to be located at different locations possibly. The sender, I'm going to call the location of the sender Q, and the receiver is going to be theta. And they're going to be located along some preference circle. It's a very standard horizontal differentiation case.

And these agents may match in the market.

So if they match they get some utility. So the sender gets this utility vs, but then has this utility for being matched with receivers that are very far away. So we have to be penalized by this distance. And the receiver is actually the same thing. So if there's a match and they get some utility, but I would rather be matched with someone who is closer to my preferences.

Not all cases produce matches. So I'm going to normalize the payoffs for not matching to zero because that could also happen.

This is just a graphical -- basically a graphical restatement of what I just said. So in this example -- so in this example there is a sender and a receiver. So the receiver could be over here at theta, which is equal here at Pi over four. The sender could be at this Q level. So that's 7Pi over four. And so the distance between the two is -- the linear distance is what I'm -- or the angle between the two is what I'm calling the distance function.

So in this case it would be a right angle, it would be Pi over two, and then we're just multiplying it by a scale of R just to have a parameter that affects both utilities at the same time. So this distance function operation utilization

that is here is just getting us the smallest difference between two locations. So that's very straightforward.

I'm going to make a couple of extra assumptions that we thought were appealing. The first one -- actually, then we look at other cases. But we start out by looking at cases where the sender has transparent motives. So everyone knows that this sender would like to match. So the goal of the dealership, for example, when I -- by clicking a banner ad or -- I know exactly what they want. They want me to go to the dealership.

And so that's going to be the case where vs is high, meaning even if I have to go to the other side of the circle, that will be a distance Pi times R, I still want the match. And then we'll look at the other cases.

From the receiver side, we're going to assume communication actually has bite. So it can be decisive. So what I mean by that is that if there is a banner ad, on average I'm not going to click it unless it says something interesting, in which case I'm interested in clicking it. So the utility extent of the receiver is not very high, but he or she -- in this case she can be persuaded otherwise.

And finally we're going to assume that the cost of acquiring information is cheap as well. If it's high, it's very intuitive, the outcome is trivial. So we're going to look at best case.

All right. The setup is also simple. The sender is going to send a message to try to induce a match. That's M. And the message is tailored through information acquisition. So the sender, before sending the message, can engage in information acquisition. And that's going to be this parameter alpha here. That's going to be between 0 and 1.

And so the way this technology is going to work is that I'm going to learn the receiver's location with probability alpha. So you can think of if the U.S. has 300 million people and the alpha is half, then with a half probability I have you in my data set, I can customize the message to you. With half probability I don't.

The receiver is going to observe alpha and the message, and based on these two pieces of information I'm going to decide whether he should match or not -- she should match or not. So this is the timing of the model. First the agents observe their own locations. Then the sender is going to choose the information level alpha. Then based on

that, sender is going to observe the receiver's location with probability alpha and is going to send a message M. The receiver observes alpha, observes M, and decides on an action and payoffs are realized. So it's a very, very simple model to set up. It's not so easy to solve as it turns out, but that's our problem.

So we're going to focus on Perfect Bayesian Equilibria, and the only thing I want to highlight here is the left-hand side. To say that the receiver is doing the following, the receiver is trying to understand where the sender is, so that's Q, based on three pieces of information. Her own location, the message she receives and the information level of the sender. So I love red cars. I see a banner for a red car. And I think, wow, that's great, that's exactly what I like. That's theta and the M is equal to theta. That's awesome.

On the other hand I think, well, is this too good to be true again because there's a high likelihood that they have data on me. So maybe actually I should think a little bit more about this. So the only thing we're doing here is making sure that the beliefs are consistent to whatever the sender is doing in equilibrium. That's it. So fairly standard.

I'll do a little bit of one focal

141

equilibrium. It turns out there's more and they're more sophisticated than this. But this is probably what people are doing in real life. We can tell the incentives for the message. So what is the message policy of the sender, right? What happens in equilibrium. And it's going to be the following: If I'm uninformed, I don't know anything about this receiver, I should just tell the truth.

So if I'm selling red cars and I'm going to show a banner ad, I have no information about this person, I might as well say I have a red car because if all goes well then this person will visit this and, guess what, I have lots of red cars and they'll find something that they like. So I might as well tell the truth.

If I'm informed, on the other hand, I'll pick some message in some set -- and I'm calling this critical set, so we'll see a theta. So I'm in different among messages as long as they convert to consumers. So maybe I know this consumer loves red cars and maybe I'll say that. Of course, cars are much more complicated than color. So, you know, I can instead -- could also present an orange car or a car that has a trim that is similar to the one that this consumer is looking for if I know that also does the

142

trick. So I may be also okay with that message.

And so we get into this general -- very general optimal communication policy, which is with one minus all the probability, I'm uninformed and so I'm just going to have a mass point here at two. I'm just going to reveal my type. Without the probability, I could have any density function here. So I could have any function on the message that depends on my location, the location of the receiver, and my information level.

All right. So this is the first result for the paper, this central result. It's this letter that we're labeling as willful ignorance and says the following: The level of information acquisition associated with the sender's first best payoff is given by this expression, this alpha bar. Don't worry about right now the expression there. The important part is that this number is always between 0 and 1.

So what's happening here is that the sender is facing a credibility tradeoff. On one hand, I would love better information because I can use that to persuade the receiver. On the other hand, if I learn too much the receiver starts understanding that the message has most likely been tailored to appear persuasive, and so there's going to be a cap -- a

143

maximum that I can learn about the receiver before credibility breaking down.

Okay. So that's the first result of the paper, is just exploring and uncovering this tradeoff between credibility and information acquisition.

In this paper, we can actually change that, but what's going to happen here -- and the idea that's happening here with this particular equilibrium is that I'm learning just as much as I can to still make it worthwhile, this click on this banner. Right? If I learn a little too much then no one will believe my plan.

This is the -- this is how this message is implemented. So now we have the preference circle here again. And here I have a receiver at Pi over two, so just on this dot over here. And maybe the sender is over here. So it's maybe Pi over four, it's, you know, nearby.

And so what's happening is the following: If the sender is uninformed, there is a matchpoint here. He's just revealing his location. That's fine. And maybe I think, oh, that's great, that's worthwhile. So an orange car is not exactly what I wanted, but it's worth investigating.

On the other hand, if the sender is

144

informed, what he's going to do is he's going to mix around this blue line because that's the density for the informed sender. So what's happening there is the following: First of all, with the highest likelihood, this sender is going to say I have the red car that you're looking for. That's the most likely case.

But then that becomes too obvious for the receiver. So the sender has to become a little more sophisticated, and sometimes choose things that are similar to what I like, but not exactly what I like, otherwise too conspicuous. And so the sender now has an incentive to mix messages a little bit.

The center could also be here at Q prime. So that's very far. That's a terrible deal for me. It's a white car. I hate -- I'm sorry, just colors. So a car that I don't like in which case I'm not interested in that particular model. All right?

So what's happening here is if you get attractive news, that could be good or bad. But if you get sort of unattractive news, you're sure that that's bad for sure. Okay? Because bad news is bad news, good news, who knows?

All right. I'm going to skip this and skip to the welfare analysis a little bit. So here what they've done also is sort of flipped the problem, and instead of thinking of information acquisition from the seller's perspective, we also resolved the model with the receiver choosing the level of information. So the receiver is now choosing the amount of privacy.

So that's another way to think of alpha is, well, if I don't do anything, the sender will have all the information that they want. But I may be -- I may want to shade my type to anonymous browsing or whatever through Ad Block in some cases. And so in which case I can decide how much information I'd like to share. Maybe there could be a market for this as well.

So what we have here on the X axis is the valuation of the sender. So that's here. And everything I told you up to now is this case of transparent motives, right? So what I'm doing here is I'm actually extending this range and I'm looking at cases where -- some other cases where the valuation of the sender is not so high. So sometimes the sender doesn't want necessarily to match with the average consumer.

And on the Y axis, I have the information level as before. In orange, I have the first best information level for the receiver, and in blue I have the first best information level for the sender.

And so I won't go through all the information that is happening here, but one of the things that stands out is we get a bang-bang solution for the receiver, meaning I either want to disclose all information or I don't want to disclose any information at all. And the intuition here is the following: If I have very niche tastes, so the S is very low, on average it's not good for a sender to communicate with all possible receives or with the average receiver.

Then I would like to share my information because I need to foster or promote communication. Right? I need to be found, right? I need this really, you know, specific comic book is very hard to find, I would love for people to be able to know that I have that very particular taste.

Same thing in the dating market. You can have a very particular thing that you like and you like your partner to also like. You're better off disclosing that thing despite the fact that now maybe people will start using that and say, oh, I also love, you know, this pretty good comic book or something like that.

The problem is after a certain threshold information is -- the communication is guaranteed.

And so after this threshold, of course you could make it more continuous, but the intuition is the same as -- once communication is ensured, then extra knowledge about my preferences is only used for persuasion. And that becomes bad for me.

On the other hand, the blue line, basically what it's saying is the sender is going to engage in as much information acquisition as he is allowed to basically by the receiver. All right?

You can actually just add the welfare measures if you think that's a good way to maximize joint welfare. I'm not sure that would be the case. As it turns out, the sender's optimal level of recognition is the same as the one that maximizes joint welfare. There is a given range in the middle that's sort of grayed out or blued, and there is just a pure transfer so there's no effects of utility. Any level is equally good on the joint sense.

All right. So we have a few other results, but this is the main thing that I like to highlight. First of all, identifying this tradeoff which is as information acquisition increases, communication loses credibility. And in the limit, suppose that these firms would like to know everything about my preferences. What would happen is I would have a

lovely time going online and just seeing everyone telling me things that I would love, right? But the problem is that I cannot believe any of it. So it gets this paradoxical effect.

So firms have it in their best interest, if we're thinking of the information part of the communication, to disclose whatever it is that they have about consumers and what's informing a particular message.

Moreover, firms are better off if you think that information collection is going to be as good as it can be, then they'll be better off engaging in partial willful ignorance about consumer preference.

And the last point that has some regulatory bite is this idea that for consumers it's not indifferent whether they can reveal or protect their data because whatever they reveal can be used to induce matches but can also be used for persuasion, even if they're strategic.

So I'll explain a little bit about what we're doing right now and then I'll conclude. So one of the things we're including is the existence of communication costs. Of course, talking to consumers in that particular case is not free. And so we can replicate the effect that as costs goes up a lot, we get exactly this thing that just the fact that you're advertising to this scale is enough. You don't have to tell me anything else. That's a credible signal.

But we also found an interesting intermediate region that is sort of counterintuitive. And what happens there is that if you're interested in communicating and paying that communication cost, you're more likely to be informed about my preferences. So you're also more likely to then try to persuade me through the content to buy, or more technically you cannot pull with attractive uninformed types anymore. So your ability to credibly convey that you're attractive decreases. So there's an intermediate region that's actually quite interesting that we're exploring right now.

About the observability assumption that I talked about in the beginning, I'd like to mention it a little bit more. So the first thing is this is a very, very standard result. If the information level is completely unobservable, so the receiver has no idea what type of information the sender has, then there's no credibility. The market breaks down or the informative part of communication breaks down. I could be naive.

So in this case, actually the sender has an

incentive to transmit alpha as best as he or she can in a credible way. The nice thing about this paper is that there's theoretical results that immediately apply to our setting that say that in the Schelling sense if -- with a certain probability, I don't observe alpha, but with a given other probability I do observe it.

Or in the next stream of literature if we pick up the van Damme and Hurkens paper, that also says something very similar. If I have a very -- if I have a noisy signal of alpha, so I sort of know what companies are doing but I'm not exactly sure, the results there are that as these signals become better, those results will be exactly our results. So if consumers have a relatively good idea of what's happening in terms of information acquisition, what it means, which is a big question, then our results will hold. We don't need to calculate those cases.

And one thing I won't talk about except mention it now is that it's actually easy to incorporate vertical competition into the same setting. It doesn't mean that we're doing it just because it's easy, but just claiming that it will be easy. And you can incorporate -- yeah. I've done that and Bagwell and Ramey have done that in different

settings.

You'll have a -- and you'll have to incorporate a holdup problem, but the good side is that consumers will actually get some utility in these markets. So it could be worth exploring.

So just a punchline that I want you to hopefully sort of provoke a little thought is that there is a tradeoff between information acquisition and credibility. Senders prefer more information because of persuasion ability, but they understand that more attractive claims now, the receivers will understand that they're more likely to be tailored. And the receiver either prefers complete privacy or complete information. Thank you very much.

(Applause.)

DR. JIN: Thank you, Pedro. Our discussion will be Upender Subramanian from UT Austin.

DR. SUBRAMANIAN: Okay. Hello, everyone. My name is Upender. I'm from UT Dallas. So thanks to Ginger, thanks to Avi and thanks to everyone else who has organized this conference. Really excited to be discussing this paper.

In the interest of time, let me just quickly get to the idea in a nutshell, what this paper is about. So you can think of many different situations.

I'm going to talk about one particular situation. There's a seller, like I'm going to say it's me, and then there's a buyer, that's you. So I'm trying to convince you to buy something. So I'm trying to convince you to buy something, I'm going to make some claims. And these are unverifiable claims. So that's kind of what sets up the cheap talk in this situation.

An interesting twist in this paper is before I make the claim to you, I can actually try to get some information about you. All right? And so that's essentially what they are studying. So what that means is I can find out what is it that you really like before I tell you what is it that I'm able to provide. And so that's essentially the setting that we have here.

And you might assume that in this kind of setting sort of the first-order effects should be as a seller, I can try to get as much information as I can about the buyer. So that would be sort of the straightforward effect.

The punchline of the paper or what's the interesting effect on the paper is that that's not always true. And why is that not always true? The reason is that the more I know about you, the more I know about you, the less you are going to believe what

I say. All right? And the important thing here is of course you must know that I know about you, right. And that's kind of the observability assumption that Pedro was talking about.

And the underlying intuition is this: the more I know about you, the more I'm going to pander. So instead of actually telling you objective information about myself, I'm going to tell you what I think you want to hear. Okay?

Now, as the buyer, if you realize that this is what is happening, that as I get more information about you I'm just going to be pandering, maybe you think of the current election cycle as people try and figure out what voters want to hear, then what I say is actually going to be less informative. And if what I say is less informative, it's going to be less influential or less credible. And that, in a nutshell, is the main focus. That's a cultivating force and that's what you need a model to analyze. The straightforward effect is you get more information, you can make more claims, but the more strategic effect is the fact that as I get more information the credibility goes down because the receiver or the buyer understands the motivation.

And then so the main result is that I don't

want to appear or I don't want to know too much about you. If I get too much information about you, I'm actually going to lose my power over you. And therefore I don't want to get too much information. And, in particular it's important that I maintain appearances. So I have to maintain the appearance that I don't know enough about you. And in the model, of course the assumption is that it's transparent. Whatever information I collect about you is known to you and therefore appearances are maintained.

So quickly what I like about the paper, it's a novel addition to the cheap talk literature. There's actually a lot of closely connected literature. This one specifically speaks to the cheap talk literature. For those of you who are a little bit rusty, you've kind of heard this jargon before. Cheap talk means three things, right? So you must be able to easily misrepresent yourself, right?

So, for example, those of you who have not met Kanishka, I would just come up here and say I'm Kanishka. And it's equally costly for me to say I'm Kanishka -- it's equally easy for me to say I'm Upender as it is for me to say Kanishka. Right? And so that's what it means when I say that the message is cheap.

Now, that does not work at the airport security, right? So if I go to the airport security guy and say that I'm Kanishka, he's going to say, very nice to meet you, Kanishka, now show me your driver's license. Right? And so it's important that it's not verifiable, the message should not be verifiable. And finally it should not be binding, which means that in this case if I say I'm Kanishka, Ginger might later come and say why don't you do the next presentation? And I don't want to do the next presentation. So there's some commitment that I don't want to get involved in. Right?

And so cheap talk means three things, that all messages are equally cheap or equally costly, that it's not verifiable and it's not binding. And we want to make sure that finally when you go to the application, all these three things are met. And there are different literatures speaking to different situations depending on which of these assumptions I make.

So Pedro talked about some of these. If you can verify it, then it becomes a disclosure literature. If the message is not equally costly, it's a signaling literature. And if it's binding, then it becomes a mechanism design on contract

literature. And so you have these kind of very closely related literatures.

What is very novel about this paper is that it focused on something that's not being cited in the cheap talk literature, which is that the seller, for example, or the sender, persuader, can collect some information about the receiver before they engage in cheap talk. And that was really cool.

They have a really nice model. There's another test and a cheap talk model that is something that the authors had to come up with to deliver the insight. And so now they've got a mathematical formulation and I really like that. And finally, of course, it has a nice insight for sort of this big data-big brother era. Right?

And this is literature, trying to understand is it always the case that given how costs of collecting and storing information are going down, are we going to find a lot of information being collected and used? And there's kind of -- Pedro's paper as well as other papers just kind of speak to the fact there are countervailing strategic effects as to why firms might self-regulate. And so I think that's interesting from that point of view.

Having said that, I had some suggestions,

157

mainly trying to kind of push this paper into the practical domain and trying to see how it might be relevant or where it might be relevant, and maybe some ideas that would strengthen the paper. Sometimes in discussion, sometimes in the actual model.

So the first thing is as you might remember, position on a circle, so technically it means that we are looking at what are known as horizontal preferences, meaning some of us might like red cars, some of us might like white cars, and so we are not all the same. Right? So that's what horizontal means.

I guess I thought initially when I was reading the paper the motivating examples were actually about best restaurant in town, and typically many product claims are of this nature, that I'm the best in town, I'm the best game in this particular, you know, something. And so that's what you would call as vertical. So it will be interesting to at least have some discussion or maybe an extension which talks about what happens when you have product claims of a vertical nature, does that still hold, when they might hold.

More specifically, the key assumption in the paper when it comes to horizontal is to say that if I

158

make a claim that I'm good at making red cars or I have a red car, it automatically means that I suck at providing any other color of car. Okay? And that's kind of the underlying forces for some of the results. At least that's what I think is the underlying force. It's for the authors to clarify whether the results will also survive, for example, if claims are neutral. Right?

So you can imagine that there are some product spaces where making one claim doesn't automatically rule you out from serving other customer segments. Right? So you might say, for example, soy milk will appeal to people for different reasons. Some people look for health, some people look for taste. Just because you make a claim that soy milk is healthy doesn't automatically rule you out in terms of taste. And so it will be interesting to know if the same results would also extend to the case where claims are neutral.

Coming to the observability assumption, right? So basically for local (indiscernible) how do you know, right, if I'm the seller and you're the buyer, how do you know what I know about you? Right? How is it exactly that you get to know that?

A couple of different things that the

159

authors discuss in the paper, so one example is, for example, Google collects information about you. Maybe they're going to tell you that we collect such and such information.

Now, in practice, of course -- and I think this might have been discussed in yesterday's forum as well -- just disclosing what information is collected may not really communicate to people what firms are able to do with that information. And within the model, you need to know exactly what the firm is able to do, the position with which they are able to do, but at least in practice may not necessarily happen.

And I think that's an important theoretical question. I mean, it's a difficult theoretical question to address, but I think as we're taking this cheap talk literature into some of these domains, I think it becomes interesting to understand how might firms actually use current mechanisms to change beliefs.

From a regulator point of view, it also throws up this question like we were talking about yesterday, maybe firms actually want to, for example, work with the FTC or other people to make disclosures about how precisely they can use this information public. Right? This is actually in the interest of

160

the firm. So one of the key implications of the current analysis is to say if firms have a vested interested to make it very precisely clear how they can use this information. Right? They don't want to hide that. By making that public, it actually acts as a commitment device and then that can actually help firms.

And from a regulation perspective, we can actually contrast this with a different type of regulation. You can either say you have to truthfully disclose what position you have, or you have to be truthful in the claims you make. And these are very two different types of regulation. The authors don't currently look at that, but I think they can have some nice implications by saying actually one kind of regulation might work well, or from more of a market welfare point of view whereas the other might not. So that I thought was interesting.

The other question is also how is this information being collected. Right? So if I'm a sales guy, if I'm trying to sell you a car and then I come and ask you what do you like, usually if you're not as naive as Pedro, you would kind of sit back and say, well, whatever I'm going to say is going to be used against me. Then you can become more strategic.

161

So in sort of the online setting, this goes to kind of ad blocking or covering your tracks. You know that people are tracking you. How does that affect your privacy concern?

And here I think an interesting result would be that allowing for people to use ad blockers might kind of be a blessing in disguise because that also regulates how much information is available to the firm. Firms may not be able to commit to how much they can collect, but we are allowing people to cover their tracks. Maybe it sets up a healthy equilibrium. Right? So that's kind of another interesting direction to look at.

Finally, I think it's important to understand whether talk is really cheap or what exact context does this apply to. As I said, as you utilize each of the assumptions in the cheap talk model, you can get into different domains.

For example, whenever there is asymmetric information, meaning that I, as a seller, know more than the buyer, then there are many standard remedies. Right? So if you want to be careful about do these remedies apply here, if they apply does cheap talk really have bite.

So, for example, a seller can back up claims

162

with a satisfaction guarantee or the fact that I have custom information about you, I'm not only going to tailor the ad I show you but I could also give you a more specific offer. I could tailor the price. And sometimes that can act as a signal of what information I have.

And so we want to kind of understand in what situations might cheap talk be sort of a fire starter problem. Again, we spoke a little about -- me and Pedro spoke a little bit about it yesterday. And so I think in markets where you can argue that there are significant holdup costs or surge costs, then meaning once you click, once you visit a dealer, the cost of visiting another dealer is too costly. That would be what I would call as a holdup problem or a surge cost problem. So markets where this would a significant problem, then I think cheap talk would apply and these results would really apply.

And so in the interest of time, let me just stop here. And if you have more questions, Pedro can handle them. Thank you.

(Applause.)

DR. JIN: Thank you. We'll take a few questions.

DR. GARDETE: Maybe everyone is enlightened

163

and there's no --

DR. JIN: Actually, I have a question --

DR. GARDETE: Oh, okay. Please.

DR. JIN: -- if you don't mind. I would just use this microphone. So your model, my understanding is there's no price. So my question is what if you introduced a price, and if the seller knows my willingness to pay, the price would be used against me, for example, and how that sort of changed your model.

DR. GARDETE: I think that's a great question. We wanted to keep the model relatively generic because in some -- you know, buying a car, there's a negotiation. If there is a posted price, there's another posted price. But, of course, if I'm buying a car, again, I'm naive, so I can be discriminated against in a good way for the seller and I may be convinced to pay more.

So there are situations where different -- you know, a seller, if he has different information about different consumers, he may be able to apply differential prices. So that's a good question.

So the idea -- the intuition isn't following: So we have a model that we can introduce that, but the intuition isn't following. On top of

164

being able to persuade this consumer, now the seller can also use that information to inform price. And so what happens is that this requirement becomes even more stringent in the sense that I can learn even less about the consumer because the consumer understands that, well, if it's a red car and they know my information on top of it, I will suffer an even higher holdup problem when I do visit the seller.

So, you know, we haven't done that, but we can explore that further. So the tradeoff being communication credibility then becomes more accentuated.

AUDIENCE: So the intuition -- I guess the takeaway if we add competition to this, is the firms are less likely to acquire information. Is that right?

DR. GARDETE: I'm not sure. It's very complicated. So it depends a little bit on what you assume these firms know about each other. So you can have -- actually, it's called a little bit the number of combinations. So it's hard to tell exactly what may happen. Can you give me your intuition of why you think that would happen?

AUDIENCE: Sure. So why exactly -- I don't know if I can explain that in -- but it seems like --

so if there's a bunch of firms out there, I very much want to -- I really need this match. And there's going to be some firm that's really close to the position of the buyer. If I can commit to not having any information, then my message is going to be most credible. And if I know there's a bunch of other people out there making similar statements, I'm going to be competing on credibility basically.

DR. GARDETE: Right.

AUDIENCE: And let me just add, it seems like there would be an interesting joint paper between the previous paper and this one in terms of collective information -- or collective reputation for not collecting information.

DR. GARDETE: All right. Here we go. Thank you. The matchmaking. So that's true, except now I know that there is a firm out there that has -- you know, is likely to have a great product. And so it can either become -- it depends a little bit of how we model it. It can even become more credible if I say, oh, I have exactly that product. I can imitate that firm as well.

The other thing that I've been a little concerned with, and it's not clear as well, is could we get into a slippery slope in which, you know, I

need more information to compete, and so given that there's another firm that already has some consumer information, I would like to compete with this firm. And so to improve my chances, I should even gather a little more information.

So can we get to a situation where it's sort of all stuck in the corner -- in a bad corner in terms of decisions. As it turns out, probably these outcomes depend on very sort of fine assumptions. So it's hard to think about these things sometimes up front. But it's interesting, too. I mean, I think that's part of the theory, in part, to think of, okay, what would happen now if we shut this off or we turn that on. And so that's the real --

AUDIENCE: Thank you.

DR. GARDETE: That's interesting. I hadn't thought about that.

Yes?

DR. COUGHLAN: I think if you put in competition, you have to start thinking carefully about the nature of demand and buyers in the market. So think about an example where all of your business is request for proposals, it's bid business. Okay?

DR. GARDETE: Mm-hmm.

DR. COUGHLAN: Then every customer is a single

customer, and the existence of competition is a pretty clear extension in some sense. But let's suppose that you're on your circle model and that in our world you're going to have to declare a price -- I'll call it a price -- or a characteristic, whatever if your characteristic that determines demand. Here's the problem. If the competitor, you're sitting here at 2:00, your competitor is at 5:00, right, and you'd like to get the customer who's at 4:30, but in order to get the customers who's at 4:30 you have to set such a low price or such a degree of redness, or whatever it is, that you then lose all the surplus you can get from the people close to you. Right?

So I think then you're -- you're in an interesting world where the specification of the nature and demand and the nature of your model, I'm not even bringing in dynamics and the revelation of type for the future.

DR. GARDETE: Right, right.

DR. COUGHLAN: But there's a ton of possibilities.

DR. GARDETE: I think -- you know, the nice project will be -- because it's significant enough, but we will have to introduce prices so it will be a different analysis in part.

DR. COUGHLAN: But price is isomorphic with red in some sense, is it not?

DR. GARDETE: It depends a little bit if you want to model in a holdup problem or not. So -- so it depends a little bit on sort of the strategy. Yeah, it's interesting enough, but first we had to take this step of being able to solve it if we have this model, then we can get there. Thank you.

AUDIENCE: This is just a quick thought, Pedro. But it seems that if you allow for many firms, competition, there could be a so-called adverse selection program that jumps in. If another firm knows that -- other firms have more information about this ad opportunity, let's say. And then I might wonder that the observations for which ads were not served actually are the worst ones.

DR. GARDETE: Right, yes, yes.

AUDIENCE: And then that might make me afraid about this rating. So then false observation, things can happen to the monitoring.

DR. GARDETE: I agree. So that's what I was trying to say a little bit. For the competition assumption, it's crucial -- the crucial assumption is to understand whether the senders know the locations of the other senders. That turns out to be very important because I don't know the end locations of the other senders then, you know, it's sort of an independent problem. But if I do know where the others are located, I may not acquire much information, get credibility, but now I'll imitate a lot of people who do know a lot about consumers. So it does become a very complex world, but we'll get there. So that's another aspect.

All right. Thank you very much for your time.

(Applause.)

DR. JIN: That will conclude our sessions in the morning. We have lunch available for you just out of this door. We request you just quickly grab the lunch and come back because we have a very interesting lunch panel. We'll start at 12:30. Thank you.

LUNCH PANEL: CAN MARKETING GO TOO FAR?

DR. JIN: Hello? We're going to start the panel soon. If you can sit down, that will be great.

Hello? Thank you.

Thank you. We have a proactive name for the lunch panel, which is Can Marketing Go Too Far? We'll figure out the answer in an hour. So, Andrew Stivers will be the moderator of this panel.

DR. STIVERS: Thank you, Ginger.

So good afternoon. I'm Andrew Stivers. I am the Deputy for Consumer Protection in the Bureau of Economics, so I serve under Ginger. So if you need to step out -- let me cut to the chase -- the answer is yes, at least from the perspective of the FTC. But I think we're going to take the opportunity here to hear from researchers across a pretty broad range of issues that are relevant to the FTC. And these are going to include information disclosure, privacy, behavioral choice, and social media.

Let me just briefly introduce our panelists. I don't want to take up too much of the time because there are more interesting things to talk about. But if you're interested, all of the biographies of our speakers are up online. So please feel free to look them up.

So immediately to my left is Jan Pappalardo. She is the head of the Division of Consumer Protection Economists at the Federal Trade Commission in the Bureau of Economics. And then I have Eric Johnson, who is a professor at the Columbia Business School, Columbia University. Next to him is Dina Mayzlin at the University of Southern California Marshall School of business; and then finally Avi Goldfarb, professor of marketing at the University of Toronto Rotman School of Management.

So we will have Jan start us off with some -- a little bit more background, a little bit more granular detail than what Ginger gave us this morning to kind of set the stage, and then each of the other researchers will present about ten minutes of their take on the research of interest. And then we'll open it up to some questions after. I certainly have some discussions -- excuse me, some questions, but I suspect that all of you will have interesting questions as well. So we will have a nice little discussion right at the end.

So without belaboring the point anymore, Jan.

DR. PAPPALARDO: Well, thank you, Andrew. It's a pleasure to be here today to be part of this

wonderful conference. Before I say anything of any consequence, I begin with a disclaimer. The views expressed today are my own and do not necessarily reflect the views of anybody else at the Federal Trade Commission, that said.

So I wanted to give you some overview of the role of consumer protection economics and marketing research at the Federal Trade Commission. I'll give you a little background on my perspective. I've been here for 30 years, came straight out of graduate school. And talk about some puzzling recent findings about the rare use of consumer research by the Federal Government to improve information remedies, and also talk about some challenges and opportunities for marketing researchers going ahead.

So my perspective. Consumer protection economics is really a relatively new kid on the block, young relative to antitrust. The Division at the Federal Trade Commission was launched in the mid 1970s. We borrow from many fields in economics, and also have borrowed quite heavily from marketing research through the years.

The Division blends research skills from consumer research with traditional economics, and I have to say that I'm really excited to see so many people interested in our area because I think there's a lot of room for collaboration going on. We're really eager to learn from you, and it's great that you're here today.

There's a really rich history of collaboration between marketing researchers and folks at the Federal Trade Commission. And if you have not seen it, I would recommend a series of essays that were published in the Journal of Public Policy and Marketing in 2014, and there's a lookback by many people in the marketing field about their time at the FTC and their experiences here.

We use a lot of research techniques that we have borrowed from marketing researchers over the years. One example is using controlled quantitative copy test techniques to try to understand how consumers comprehend marketing messages. Classic cases where there's actually quite a bit of literature in the academic realm is a classic case, FTC v. Kraft and FTC v. Stouffer Foods.

We worked with and learned from consumer research and marketing researchers and have used -- the agency has relied on marketing researchers and a lot of their cases using consumer surveys. An example of that is FTC v. Dolby, evaluating customer success,

and FTC v. TransUnion, evaluating consumer attitudes toward the use of information from credit files to compile marketing lists.

We've used empirical analysis of consumer behavior increasingly in our cases. And increasingly it's become more sophisticated with more availability of granular data and bigger data sets about what firms are doing and the overall marketplace. An example of that is a finite mixture modeling piece that was recently made public in RIO. It was worked on by Devesh Raval to identify types of content providers largely responsible for cramming in the T-Mobile and AT&T case.

And one thing I would mention is that a lot of our work is private, right? So you see the tip of the iceberg of what the FTC does. There's quite a bit of work that's done behind the scenes and investigations that incorporates a lot of very -- demand analysis, consumer research, really quite a range of things. And I wish we could bring it all to your attention, but the nature of the beast is that the publicly available cases are the ones where you get a sense of what's going on behind the scenes.

We've done content analysis. Many, many years ago, I was very interested in trying to

understand how advertising regulation actually affected the types of health messages that firms gave to consumers in marketing. And I was lucky enough to have been at a marketing conference and tell some folks that this was something I was interested in. And they said, oh, if you're interested in content analysis, you should pair up with Debra Ringold to do some research in that area because she had specialized in content analysis.

We did research that was later published in the Journal of Public Policy and Marketing, and then I worked with another colleague to do some content analysis. And that other colleague is Pauline Ippolito.

We've done surveys and experiments to study consumer fraud. I think Ginger mentioned earlier today that Keith Anderson has taken the lead on doing many surveys to try to estimate the incidents of consumer fraud in the United States and something about the characteristics of people who are likely to be fraud victims.

We've done a lab experiment. Folks have worked on trying to understand the characteristics of folks who are likely to be deceived. We've done controlled experiments to assess disclosures,

appliance energy labeling and mortgage disclosure research, and I'd like to talk a little bit about that in more detail.

The energy labeling question was one about what type of label the FTC ought to use to convey to consumers accurately what types of energy features there are on appliances. And at the time, Congress suggested that we might want to go to a star or a categorical label. At the time, we were using a label that featured kilowatt hours. So we said, well, why don't we test that.

And we worked with colleagues in the Bureau of Consumer Protection, and we did an online panel study, controlled, randomized experiment. And it was a very interesting study, because in addition to doing the star label and the kilowatt-hour-featured label, we decided to test one that featured a dollar metric.

So what were the bottom-line findings? What we found was that overly simplistic metrics, such as stars, can actually hinder consumer understanding. People seem to think that the star meant something more than the energy efficiency attribute of the product and applied to other features of the product. In the end, based on this research and public comments and analysis by FTC staff, the Commission decides to

177

go to a label that featured dollars as a key metric. We found that dollar amount metrics are meaningful, and this is intuitive, because people can use dollars to compare across all kinds of goods and services. They're trying to figure out how to optimize utilities subject to their budget constraints.

We did mortgage disclosure research because we found in cases at the Federal Trade Commission that consumers could be totally clueless about the features of their mortgages, even if they had received the federally required mortgage disclosures. And we were wondering, is there something about the disclosures themselves that could be a problem, and is there something about the disclosures themselves that could be improved to help people make better decisions.

So we did a two-part study. We used in-depth consumer interviews for the first part. We talked to recent mortgage borrowers. And we also did a quantitative randomized, controlled experiment, testing what were then the current disclosures, and good versions of the current disclosures, I might add, against a prototype developed here at the FTC.

What did we find? Well, the qualitative research was fascinating. We found that many people were unaware of or did not understand key costs or

178

features of their loans. And even worse, we found that some of the mandated terms were actually misleading to consumers. People thought a discount fee was not really what a discount fee was.

We developed a prototype disclosure; we did controlled testing. We found that people did substantially better if we created a document with the first principles of what would you want your best friend to know if your best friend was shopping for a mortgage. And we used features from consumer research to try to say what is clear, what -- how do you layer the information so the most important information is on the first page and so forth.

We got substantial improvements with the prototype versus the alternative. We found that extraneous information with additional details can confuse consumers; descriptors can be misleading; and controlled, quantitative consumer research can substantially improve disclosures and may be necessary to avoid inadvertent deception from well-meaning disclosures. So we know this. I think people have known this for 30 or 40 years. You really need to test in controlled settings consumer understanding as possible consumer behavior in field experiments.

Here's the puzzle. A recent study found

179

that government agencies rarely use consumer research in their decision-making. Fraas and Lutter found that although federal mandates to disclose information underpin a number of flagship regulatory initiatives and sundry major regulations, we've only found a very few exceptional cases where there's any evidence that the responsible regulatory agencies conducted research.

So here's a question for all of you in the marketing field. Why is consumer research not a routine part of consumer policy development? Do policymakers not recognize that well-meaning disclosures can mislead? Do policymakers understand the potential benefits of consumer research but think the cost generally does not outweigh them? And what are the costs and benefits of alternative methodologies?

A few hot research questions for you guys to think about: how to provide reliable estimates of consumers' willingness to pay in markets without market prices. This is very important for the world of privacy and data security.

How do we translate established techniques for advertising disclosure testing in traditional media to newer media? There was a discussion of that

180

yesterday and I think today as well. Very important question.

There are many opportunities to try to collaborate with folks at the FTC. In the past, we've had people work jointly on projects. We've had people come for sabbaticals. And I think it's really helpful to just talk to people at the FTC who are on staff working in your area as you develop research projects outside of the Federal Trade Commission to make sure that you understand the nuance of the policies, of the law, the regulations, and the policy questions to make sure that your hard work is as relevant as possible to the real world. And I thank you very much. Oh, I have some references if you want references.

(Applause)

DR. STIVERS: Great. Thank you, Jan. And now we have Eric.

DR. JOHNSON: Thank you. It's nice to know I can still talk IBM if I need to. I have to say one thing because Jan said it well. I've been spending the last years as a senior visiting scholar at -- sort of across town at the Consumer Financial Protection Bureau, and that's been wonderful, but that means I have to use the same disclosure. So what she said.

The best version of that is someone who adds

181

"and it's not even the opinion of the United States, but somewhere there's a country that approves of what I'm about to tell you."

What I'm going to say -- essentially three things. One is I want to introduce why regulation should take a behavioral perspective. The second thing I want to do is offer an example, a contrast example, for mortgage decision-making partly inspired by some great work that Jan just talked about that she was involved in. And, finally, I want to start with some -- stop with some observations about disclosure.

Okay. So I think now the field has matured, that we actually have some good empirically grounded models of how people behave that are departures from the standard analysis. One of these is basically models, and I'm going to think particularly of beta-delta or quasi-hyperbolic discounting of time preferences. And I'll come back to that because I think it makes all the difference in the world when you talk about mortgages.

Another example is we know a lot about risk preferences, and we know about, a lot about loss aversion. And, finally, you know, we can put a quick view of it, I'll call it limits on information processing. They're very clean models that people do

182

not necessarily think all the way down the tree. There is a notion of K-level reasoning -- we don't go to the bottom.

I'm only going to talk as an example about time preferences. The reason I think this is so critically, critically important is because if you don't include these models, you're going to end up with not only results that are wrong but that can hurt social welfare and actually hurt social welfare in a way that hurts the most vulnerable people.

And I'll illustrate that in two examples, but you can imagine just one quick thought, if I'm doing disclosure and you think that people have costs of processing information, those costs might be correlated with education or socioeconomic status. Disclosure might actually be harmful or at least not as helpful for people who are not as well off.

Okay, so, let me give you my favorite example, and this is a paper that's in press in the Journal of Marketing Research with Steve Atlas, who was a Ph.D. student at Columbia, and you might know who this guy, John Payne, is. So, there are two kind of mortgages in my world. Not only am I going to tell you about our toy model, all we have is a toy model. But essentially imagine the two mortgages. What is

183

commonly called a 2/28. For two years, you get a great rate, and after that two years, you have a terrible rate. Okay?

Typically this has -- and this is important -- no money down. So, it's a great rate, no money down, and you get to move into the house immediately. The other is, of course, the classic old, boring, 30-year fixed-rate mortgage. Now, if you think about this from a principal agent problem, this was a beautiful device. Okay, people who were creditworthy who get 30-year mortgages did, but there are people out there who know they're going to have good credit ratings in two years.

So what they're going to do is take the 2/28. And if I'm not good, I won't buy a mortgage. Sounds like a beautiful separating equilibrium, right? Now, what else could the 2/28 mortgage be? What is the kind of person who it might appeal to? Imagine you believe in present-bias or hyperbolic discounting. What happens in this analysis is very simple. The 2/28 becomes a present-bias magnet.

And I don't have to tell you how this story ends, you know, not well. And if you've seen a couple of recent movies, you might know. We did an analysis of this where we essentially did two ways of data

184

collection. One is we actually managed to get questions about loss aversion and time preference in a nationally representative sample, which actually turns out to be done by the industry together, three hours' worth of survey data about people's finances.

And the other thing is we actually did our own survey using DEEP, which is a technique that Olivier Toubia and a bunch of us have developed, which gives you basically -- it can give you time preferences in a beta delta model in about eight questions or actually parameters from a cumulative prospect theory model in about eight questions. So it's way cool, I think.

And basically our little toy model says three things. First is present-bias and impatience will make people choose adjustable 2/28 mortgages. I mean, that should be clear as an intuition. Second, if there's a shock -- and here I'm talking about a negative shock -- to house prices, because they have less money in the mortgage, they will, in fact, be more likely to be under water, okay?

And the standard analysis, if you read the press in 2008 and '09 is many, many people would walk away from such mortgages. It would be cheaper for them to move and rent and leave the balance. But our analysis, and we borrow this largely from Dellavigna and Pollet who've talked about it in labor markets, is, think about it, if I bought the mortgage because it didn't hurt me at all to get it at first, think about walking away. Walking away has huge costs. You know, many of them nonpecuniary, but, you know, I have to move, I have to find a rental, like change kid's school, et cetera, et cetera. And the benefits are delayed.

So what this suggests is actually the reverse to what -- you'll not only more likely get into the bad mortgage, but you're more likely to stay in the bad mortgage, which is essentially the analysis from labor that Dellavigna and Pollet basically did using these two data sets and lots of controls, we basically showed that present-bias leads you to get adjustable rate mortgages and keeps you from walking away.

And I just want to contrast this very pretty model, the 2/28 separates people into creditworthy and non-creditworthy to what I think is the reality, which basically became not only a magnet for people with present-bias but they were condemned to that situation over a long time. Now, I've just pointed out the observation. The first version of the Household

Affordable Refinance Program, called HARP, which put in $7 billion to try and get about 7 million people to refinance is largely considered a failure because only 2 million refinanced. So, I mean, it's consistent with this story.

Okay. Now I'm going to shift from talking about mortgages to talking a little bit about privacy and disclosure. One thing I want to point out, I think the notion that people have a utility for privacy is probably a little naive. It's what I call an assembled value. For those of you who know the term "constructed value," it's my substitute for that because assembled means I have lots of things. I want to have customized products, but I also don't want you to sell my information. And how those get thrown together is a function of how I ask the question.

We did some research that was published in the Communications of the ACM, where we essentially did the old opt-in/opt-out, which we have done with organ donation and other things. If you have people having to check in to get more mail surveys, in this case only 48 percent of people did. If they opt out not to get them, 98 percent of the people would get these surveys. So, you know, the same standard story. How you frame it is how -- what will be, quote,

unquote, revealed.

Okay, two last comments about disclosure. One is I want to remind you of the work of George Loewenstein and many other people who show that disclosure can have perverse effects. They looked at the setting where a doctor will disclose I own the lab and I'm sending you for a test.

What they find reliably is that people say, oh, he's a nice guy, he didn't have to tell me that. In fact, people are not more suspicious; they're, in fact, less suspicious. So in that particular disclosure framework -- and disclosure is much more complicated than that -- it's problematic.

The second thing I want to point out is it raises processing costs. As I said earlier, if processing costs are differentially available to different folks, and I'll use an article -- an example from Ben-Shahar and Schneider, who have a nice book called The Failure of Mandated Disclosure. Actually, that's a law review article, which is cheaper than the book and has all the good content.

But if you read the law review article, they close by what is the effect of hospital quality disclosures. Yeah, they're kind of hard to read and hard to find, but they say basically what they believe

happens is that wealthier and more educated people, in fact, find the good hospitals, go to them, and as a result what beds are left over? The ones at the not-so-good hospitals. And that disclosure actually does maybe improve consumer welfare for some people but not for everybody.

So I just want to point out disclosure gets to be kind of interesting and complicated as soon as you assume information is costly to process.

Finally, the solutions, and this is -- I'm writing a book on choice architecture, so I'm going to make a plug for this, which is lowering processing costs through choice architecture. And that's a whole other talk, so I won't talk about that much now, but I'd be glad to talk to you about that later. Thank you very much.

(Applause)

DR. STIVERS: Thank you, Eric. That's going to be followed by Dina.

DR. MAYZLIN: So, my name is Dina Mayzlin. Thank you very much for having me. So I'm going to talk about consumer welfare and regulation of social media. I'm the shortest speaker here. And I'm primarily going to -- basically I'm going to talk about two papers that I've done on this topic, and

189

then I'll talk about some other things that I haven't worked on and not that many people have worked on, but I think are sort of interesting and intriguing.

So what is social media? So some of you may think of social media as the platforms. I actually have my Facebook friends -- Eric Johnson is one of them. His picture is there. So I think of social media as -- so the medium are the consumers, okay? So instead of sort of a firm advertising on -- you know, on TV or on print, here the consumers are talking to each other.

And, so, you can, of course, think about the platforms as well. And, so, you know, usually when I give this talk I talk about the role that the firm can play in managing social media. But, of course, here, we have a slightly different perspective because -- and this actually -- I don't know, the first time I'm interacting with the Federal Government so, you know, we're more worried about perhaps consumer welfare.

So since we're worried about consumer welfare, let's think about consumers and how they use social media. So I'm going to -- you know, I usually talk about the three Cs of social media, so connection, curation, and content, where content is like the stuff you read, perhaps it could be blogs,

190

word of mouth. There's also connection where you're connecting with friends.

And curation is the idea that, well, I'm going to follow someone on Twitter because this person knows a lot about interesting things that I should be reading about. And different platforms have these kind of different uses, so I would argue that, you know, Facebook is largely about connection; Twitter is about curation; things like blogs are about content.

Okay. So why should we worry about regulation? And I have to say that this area has not been well regulated, I think it's okay to say. There has been some regulation by the FTC, but it's kind of pathetic.

(Laughter)

DR. MAYZLIN: Pathetic in a good way, in a good way, in the most positive use of the word. I mean, there is reasons why it's -- I mean, pathetic in the sense that there are some -- basically, if you don't disclose your connection there may be a fine you can pay. Very few people have been fined. And, so -- and there are reasons why it's so hard to do it, there are all these different players, you know, it's kind of a nightmare.

And it's also been good for me, my own

191

research stream, because I've been -- you know, when I was in the market in 2001, I was kind of worried about, you know, oh, the Government has cracked down on this whole review manipulation, but thankfully it has not. So...

(Laughter)

DR. MAYZLIN: So I was able to get more papers out of it. All right.

So the idea is that, you know -- so, again, I usually talk to firms about managing social interactions. And, again, the idea is that in their management of social interaction some may be legal, some may not be, you know, ethical or unethical. You may have negative impact in consumer welfare, and we'll talk about that.

And the second thing, which I haven't done research on and I don't think a lot of people have done research on, is that, you know, I think the thing that worries me a lot now as a parent and also as a researcher is what is happening with social media -- misuse of social media.

And there's sort of two things I've observed in the past year that's been a big deal. The first one is this idea of incitement of political -- and some of it has to do, you know, may have to do with

192

terrorism or kind of racial incitement. And the second is this misuse of social media by minors and the long-term consequences that can have for kids.

And, so, you know, our Government usually worries about -- the Government usually worries about, you know, the area of, you know, minors, and so I think it's kind of a big deal. It's actually, I think, one of the biggest issues that schools now face is the use of social media by children -- schools and parents.

Okay, so just kind of a few more frameworks on this. So there's usually -- I think about sort of three different roles that the firms can play in managing social interaction. So one is very passive, which is listening. And I think a lot of sort of what, you know, we've talked about, collecting information, you know, tracking; not necessarily really acting on this information is being done, so it's being done, you know, all the data is being scraped, multiple companies, and, you know, ostensibly, you know, anything that's online is scrapable, and so that's then collected.

You can also think about more kind of active roles. So one is engagement, where you try to get people to talk about your product. Another, and the most kind of aggressive one, is promote, where you try to sort, you know, get people to buy stuff through social media. And, again, some of it may be done with disclosure by the firm; and some if it may be done without disclosure. And, so, that's what I'm going to talk about.

So I have these two papers on this topic of what happens when firms try to manufacture word of mouth, so basically try to pretend, to enter the conversations but not reveal that they're -- that they are there. And, so, that can be done under -- you know, because virtual space provides you anonymity.

All right. So I have -- so, my -- you know, it's a long time ago. My job market paper was on this, and it's this idea of promotional chat. So, it's this idea that, you know, we saw -- you know, I saw back in the day that people started to talk about, you know, CDs, music, movies on online forums, and there was a case of a singer that basically her representative is one in these online forums and pushed -- pretended to be kids.

This is one of the cases that Ginger talked about, right, with the Sony case. It was basically that, just, you know, a few years later. And, so, you know, I started -- when I saw a case like that back in

-- I think it was about '99, 2000, started to wonder about, you know, so what does that mean. So we basically have advertising where we can no longer tell if it's real word of mouth or advertising content or paid content.

And, so, you know, so I think the interesting thing was to look at the equilibriums. So if you think about the consumers now know that this is going on, so they know that these -- in my model they're competing firms doing this, do they -- you know, does it still work? Can you still be persuaded if you think that people have these bad incentives, does it just fall apart?

And, so, we find that in equilibrium you still have an informed equilibrium, so it basically kind of -- what happens is that this basically adds noise, so sometimes you're going to be making wrong decisions because some of the -- some guys are getting messages that are false, that are just promotion. And, also, an interesting thing is that the worst product is going to be doing more of this.

So, but despite this, because of real word of mouth, there is kind of truth-telling that happens in equilibrium. So, on average, you're okay. But if you actually saw how much firms promote, you would see

that the worst firm would promote more, would invest more into this.

And, then, you know, if you think about welfare, I think by the end the paper was published, there wasn't much welfare left, but initially there was more welfare. They have to make it short, so there was a bigger section of welfare and you could look at consumer welfare, so you can kind of look at -- you know, so, of course, if -- so basically one of the results is that as it becomes more costly to do this, there will be less kind of fake reviews in equilibrium, and so you're going to have, you know, more consumer welfare.

Also, the extent of the real chat matters. So if, you know, there's not enough, then there's going to be a lot of noise and signals, so you're going to be making, you know, kind of bad decisions all the time.

But, I mean, so I think -- so one thing I want to highlight is this idea that, you know, of course we don't want there to be bad reviews out there or fake reviews out because people are going to be making wrong choices.

But I think even a more important kind of, you know, negative consequence that could happen is

that the whole thing could fall apart, right? So if these things become so spammed that nobody wants to go there, which I would argue would happen to IRC -- I don't know if you guys remember IRC back from my -- nobody remembers, okay. It was this online channel kind of chatrooms way back in the day, basically they got spammed and disappeared. How many of you remember IRC? All right, all right.

Okay, yes, yeah. So those things became much less popular, and so, you know, I think you think in terms of welfare, you could think about the noise added to the -- but I also think, well, you know, is this something that will destroy online forums, online communities. And I think by now we're sort of -- you know, I feel like it hasn't destroyed. You know, we can say with more confidence that it's not going to destroy it, but it's definitely going to add noise.

Okay. And then another paper, kind of a more recent paper I have with Judy Chevalier at Yale and Yaniv Dover at Hebrew University in Jerusalem is actually an empirical paper of the same topic. You know, it took us a while to write the followup empirical paper, and the reason is that you kind of, you know, couldn't -- I don't know, I and probably other researchers sort of couldn't think of a way to

197

really study this phenomenon because sort of by definition you're saying you can't tell it apart, right?

So part of the kind of setup of that model was you don't know if it's coming from a consumer, it's coming from an interested party. So if you don't know by definition, then how do you study this thing, you know, it just becomes sort of -- completely unobservable by definition.

And, so, this paper, what it does is it exploits a variation in platform design. So as -- you know, as the space has evolved, Tripadvisor and Expedia have very different design features. And one of the design features is that Tripadvisor allows everybody to post a review; and Expedia verifies the authenticity of their reviewers. So they basically -- they just make sure you booked the hotel through Expedia. Okay, and if you didn't, they're not going to post your review there.

And, so, we use that, along with variation in kind of organizational structure. So some hotels have small owners; some hotels have large owners; some hotels happen to be right next to a competitor that is a small owner, large owner, independent or chain. And we have sort of -- basically assumptions on, you know,

198

who has -- using literature, kind of economics literature, organizational structure literature on, you know, who has more of incentive to fake.

And, so, we know which hotels have more incentive to fake; we know who is collocated next to whom; and then we also know kind of -- we look at the difference in the -- so we don't look at each review because by definition we don't know how to do it; we can't tell it apart. But what we can do is we can look at the difference in the distribution of reviews. So we basically look at the same hotel; look at the distribution of reviews it has on Tripadvisor; and compare it to the distribution of reviews on Expedia.

And we see the extent to which they differ. So, for example, what we would hypothesize is when you're next to this very aggressive competitor that's small and independent, then you're going to start to see -- you're going to pop -- you're going to have more negative reviews on Tripadvisor relative to Expedia because it's easier to fake, and Expedia is a bit harder to fake. And, so, that's our methodology.

So we basically, you know, use this kind of, you know, where we use this to kind of -- we propose as a mechanism to detect review -- to detect some sort of fakery, these differences in distributions, and

199

then we see the extent to which we see fakery. And, so, let me give you kind of an example of what we found that summarize our results.

So you can compare -- so, again, we don't know -- we can just tell the difference. We don't know the absolute level of fakery, but we can compare Hotel A that's a branded chain and a large owner, so sort of less likely to fake. Hotel B is an independent and small owner that we think is more likely to fake. And what we see is in the data, that Hotel B will have seven more five-star reviews on Tripadvisor, and the average number of five-star reviews on Tripadvisor is 37.

Okay. So, and this is sort of like a reasonable result, I think, because it's not like overwhelming, right? Like it doesn't kill it. But at the same time, you know, it seems pretty big. You know, it's -- so it's adding noise to the signal.

Then if you look at I think a more interesting result is this faking negative reviews for a competitor, which seems even kind of more, you know, aggressive. So, if you can look at -- if you look at Hotel C that's located next to again this kind of, you know, a bigger, less aggressive faker, branded chain, a large owner, versus Hotel D that's located next to a

200

more aggressive faker, you -- we find that Hotel D has six more one- and two-star reviews on Tripadvisor. And the average number of one- and two-star reviews on Tripadvisor was 30.

Okay, so it's a pretty significant amount of, you know, negative reviewing as well. But, again, it doesn't kill it, but it seems like, yeah, kind of a big deal.

So, all right. And I don't want to suggest, you know, there's been some other papers. I know that Michael Luca has a paper that's either forthcoming or about to be forthcoming on also fake reviews, or it already came out. So there's a few other kind of recent papers on this. But I just, you know, just want to talk about my own stuff.

So the last thing I want to talk about that -- actually wanted to show you guys a video, but the reason I couldn't show it is that the Federal Government allows, like, sharing sites to be accessed through their computers, so I couldn't share -- I didn't think of that.

So -- but the video was -- okay, so let me talk about the first point. The first one first, and it's really just kind of speculative, like I don't have a model to show this and, you know, but the concern is that -- so I saw this -- I mean, I think you see this in elections, right, during the election season.

You just see these kind of, you know, crazy conspiracy theories, you know, kind of spinning out of control, and you notice that, you know, people seem to be just in different worlds, you know, like the -- depending on whatever your political affiliation is, you're just getting different news, and news just seems to get very, very extreme. You know, news/opinions/conspiracy theories.

And, so, what is going on? So one hypothesis is that if you have homophily in social networks, you kind of have this amplification effect of social media. And it seems like, you know, extreme content seems to propagate. And I -- you know, I think it would be an interesting thing to show -- to show that in a model.

And I think it's a big deal. You know, it's a big deal if you think about, you know, our role in the Middle East or perhaps, you know, what is said about -- you know, think about the -- but you also think about the kind of local domestic policy, you know, the fact that people get so much of their news from social media, and they seem to be getting kind of

very, you know, twisted version of the truth, you know, has a big effect on, you know, our country and how elections run and how, you know, public and foreign policy develops.

The second point is a point about use of social media by minors. And, so, I mean, how many of you have kids who are, like, between the ages of, you know, 11 and 18? Okay, so a few of you. So, you know, it turns out that this is kind of a big deal. And, you know, and I blame Snapchat, one of the platforms.

And, so, what happens in -- these social platforms are very popular among very young children, so starting, I would say, with the age of 10 or 11, kids get their smartphones; they get these apps; and there's very little monitoring by parents. You know, they're basically on their own.

But the problem with that is, you know, and you could imagine that kids of that age don't -- you know, they basically don't realize the implications of their behavior. They gauge their behavior for themselves or for their friends. And, then, you know, there are like these sort of things that spin out of control.

The other thing that I think is very

troubling is that the electronic footprint doesn't get destroyed. So your fear -- you know, this interview I was going to show you was an 11-year-old saying from my kid's school, you know, I interviewed them for my class, saying how sexting is very popular in sixth grade.

And, so, these are basically 11-year-olds, you know, kind of texting pictures of themselves, you know, naked pictures of themselves to boys. You know, it's usually -- it's usually girls to boys to kind of impress them. But that stuff, you know, basically, as soon as the boy gets it, he forwards it to everyone else in his circle.

And, so -- and I think part of the reason that happens is that sort of the normal pressures of growing up and trying to kind of fit in and the fact that social media is about connections, but part of it is also there's kind of false sense of -- you know, Snapchats, the stuff is supposed to disappear after a few seconds, but you can take a screenshot, right?

So you don't -- you know, as soon as you get that picture, take a screenshot, and so it doesn't quite disappear. So, you know, first of all, they're too young to probably understand, but also, they don't quite understand the technology.

And, so, I'm not sure exactly what the solution is. I mean, one solution is just to be stricter about not allowing minors to use it, you know. I know my kid's school basically they this year outlawed the use of smartphones during school hours. But I have to tell you, I talked to the principal or the superintendent of the district, and he said it's, like, one of the biggest issues they face -- sexting, cyber bullying, and bomb threats. So social -- there was -- my district had three or four bomb threats last year using the site Yik Yak, which provides anonymity. You can post anonymously and just like -- they just went out of control. So I'll just leave you there.

(Applause)

DR. STIVERS: Thank you, Dina. And we will finish up with Avi. And, unfortunately, we may be running out of time, so this may turn into more of a lightning round than a panel discussion, but maybe we can grab a little bit of time. So we'll see.

DR. GOLDFARB: Okay, so, before I start, I should say that all -- pretty much all these ideas, including many of the slides, were developed in collaboration with Catherine, so actually there's a couple papers that are hers and not mine that I will be citing. Okay.

205

So what's privacy? Privacy is the right to be left alone -- left alone and the right to no unauthorized intrusion. This is a hundred-year-old definition, and in the law, up until the last few years, privacy was something different. Privacy was a public versus private life distinction. Public figures had the expectation of having their picture taken in certain places, and private -- you know, private figures, if you were not a public figure, you didn't have to worry about that kind of thing, and there was a distinction in the law.

Or there was a sense of privacy and security, whether you're going to be wire-tapped or whether -- and it's very much about government surveillance of individuals, which is still there, but privacy is now a business issue as well.

And, so, what's happened to make privacy a business issue? It's that data is now key to innovation in lots of industry. So, you know, I quoted a couple of leading economists on this, one who tends to be very much thinking about the future, Erik Brynjolfsson; another who is an historian, thinks about the past, but also says, hey, if we look at digital age, data is fundamental and it seems to be changing things in a deep way in terms of innovation.

206

And it turns out that the use of data requires data. And that means that privacy regulation, if you think about it, is about explicitly restricting the collection and use of data. Privacy regulation is about restricting data flows. And if we need data for innovation, this could be difficult.

But it turns out that consumers and governments, as, you know, we've heard the word "privacy" a lot today, and I heard it a bunch yesterday, are concerned with threats to privacy. So companies can use data to harm consumers by charging higher prices or denying service. There's also this big element that it's hard to really define, even when you push people, that it's creepy or repugnant that companies know more about their life than they do.

So, as a consequence, we've seen some regulatory attention, sectoral in the U.S. and more general in Europe and to some extent in Canada. Okay.

But then when you look at what people do, and this is related to what Dina just said, is maybe people don't care as much about privacy as they seem to. And, so, how do we reconcile these issues and how do we think about privacy when we acknowledge that maybe people don't care in certain situations or people are revealing a lot about themselves, and then

207

we look at the consequences of a lot of the regulations that we do have; they restrict innovation, and they hurt outcomes in the context of health.

So underlying all this, I think, is an idea that privacy and openness are both positive values. So we want privacy, but we want openness. And in particular in an innovation world, we think about how do we facilitate innovation, how do we foster innovation, openness is fundamental to that.

But privacy and openness are opposites. And, so, we have two positive values that in many ways conflict. So this suggests we're going to have some kind of tradeoff between privacy protection and innovation. And, so, this is pretty bleak from the point of view of thinking about privacy regulations, and consumers seem to care about this, or at least they say they care about this, but we -- are we really willing to hamper our economy in some way in terms of innovation?

And, so, there's a question of, you know, maybe we should just have a free market and why regulate this thing at all. So I'm going to start with the premise that consumers actually do care. So consumers do react negatively to some kinds -- not all kinds but some kinds of privacy-intrusive advertising.

208

Catherine and I showed that in a paper about five years ago, and we've seen more evidence of this, is that if you violate privacy in the wrong way as a firm consumers get really angry at you, or at least they stop buying from you and they behave differently.

Second, over time, consumers are becoming more reluctant to share data. So if you fix the context of sharing data, in this case it's do you give your income in a survey, people over time are becoming less likely to share. So what's changed is that the benefits of sharing have grown so much relative to the cost. So even though people in a given setting share less, maybe there's -- the benefits to sharing and social media are sufficiently high and have grown so much that we seem to see more of these -- more sharing.

So how do we think about privacy regulation? I think the -- one privacy regulation that seemed to foster both innovation and consumer protection was the Fair Credit Reporting Act. Okay, so, this, at the end of the day, is privacy regulation. It is about how do we regulate consumer information about credit. And an important aspect of it was there was a centralized repository where consumers could go and figure out if information was accurate, and that actually helped

209

firms, too. This was sort of a really nice win/win. Consumers could figure out what firms knew about their credit; and firms could have some verification when that was wrong.

So I think one of the most useful things to think about in the context of privacy regulation is to try and figure out if there's some kind of regulatory model around clear and consistent disclosures that's like this Fair Credit Reporting Act in the context of online. And I don't have a good answer. I mean, that's just a question, okay?

But, so, now what do we do? I think we have to -- when we think about privacy policy, we think about consumer protection, but we also think about innovation. And it can't be too strict or else it's going to stifle data-driven innovation, and that's the work that Catherine and I had started working on about five years ago, or at least published five years ago. But at the same time -- we worked on it a little bit before that.

But at the same time, privacy regulation can't be too lax. And this is what we're starting to see, or else consumers will be unwilling to provide data, and again it's going to stifle data-driven innovation. And getting the balance right is going to

210

be the key challenge in the future. And, so, we might think about, to the extent that we want to think about the Fair Credit Reporting Act, is there some way we can enable disclosures and almost have more openness about privacy. Thanks.

(Applause)

DR. STIVERS: All right, so how are we doing on time, Ginger?

DR. JIN: Fifteen minutes.

DR. STIVERS: We have 15 minutes?

DR. JIN: Laura says 15 minutes.

DR. STIVERS: Oh, great, okay. Well, then, first let me thank our panelists. Let me go ahead and open it up to the audience first and see what questions we have.

AUDIENCE: Avi, when you talked to your comment on innovation, are you thinking of private innovation? Innovation is coming in firms, right, or is it public U.S.-based? What are you referring to in your context?

DR. GOLDFARB: So, I was thinking firms, but it also is related to universeness. So if you think about all the research --

AUDIENCE: Okay. My question is for these firms --

211

DR. GOLDFARB: Yeah.

AUDIENCE: -- why do you think they should disclose before they are ready for a new product or whatever?

DR. GOLDFARB: Why do I think they should disclose?

AUDIENCE: Yeah, I don't know what you are saying there.

DR. GOLDFARB: Oh, so, I think there's potential for consumer harm from use of data, in particular the fact that data is -- information is non-rival and so the firm can collect it and then the consumer might have very few reached rights on what happens to the data after it's been collected. That's a potential aspect of harm.

So at the same time, all of these regulations we have, at least the ones we've seen so far, primarily in Europe but a little bit here, are not just hurting the firms' ability to profit from data but also hurting the ability of the firms to help consumers, and in the hospital case, save lives.

So to the extent that there's some way to think through letting consumers know what's happening with the data -- okay, so I should be clear that I don't have -- I don't know what the right policy is.

212

I know the right policy isn't -- I know it's not a good policy to say you can't use data. Okay? And, so, given that in the presence of knowledge the consumers seem to care about this in certain situations, how can we make them able to make informed decisions so that we can still have innovation and the consumers are still willing to provide data to firms so that firms can better serve the consumers.

AUDIENCE: Avi, I'm just curious to know, I'm having a difficult time understanding where the market failure is that we need to have some regulation to actually correct this market failure in the innovation.

DR. GOLDFARB: Okay. So that's fair, I skipped that. So the fundamental market failure is that information is non-rival. So once the information -- once a consumer provides information -- the potential for market failure, I should say, is that information is non-rival. So once a consumer provides information to a firm, that firm can share that information and keep it. And it doesn't need to tell the consumer about that.

So that can lead to a variety of interrelated market failures. So one, for example, is that we can get complete unraveling of markets, so that consumers are unwilling to buy from firms because they're afraid that that firm is going to share information about their preferences with other firms.

AUDIENCE: But won't the firm then realize that and then just start to close that --

DR. GOLDFARB: But if there's no way to commit -- so if there's no credible way to commit -- because the information is non-rival, then it gets -- so this is -- there was a handful of papers, Curtis Taylor has a paper on this and Alessandro Acquisti and others, in a paper on this that came out around the same time showing that markets can unravel, and we have some sense of that.

DR. STIVERS: And I should be clear, given this is a panel, everyone else up here is welcome to jump in to answer.

DR. JOHNSON: Just one small point, which is that assumes people have a known preference, they understand the problem. And if you see that they change their preferences depending upon whether we checked the box or not, that would make that assumption questionable.

DR. MAYZLIN: So I'm going to add this. It's not actually my research, but Alessandro Acquisti has this really cool paper that shows that even, you

know, just having someone's image is enough to connect -- you know, to connect your data to your, you know, birthday, and once you know your birthday and hometown, I can get your Social Security number, and once I can get your Social Security number, I can do all kinds of things.

And, so, like this information is very basic that's shared on Facebook, which people like to do because they like to get happy birthday wishes. The birthday and your hometown actually is, you know, incredibly useful information if you want to know someone's Social Security number.

DR. SUBRAMANIAN: My question is for Eric. Eric, you mentioned that, you know, we should not specify your utility function, for privacy that it's an assembled construct. I think a big question is to -- from the researcher's perspective is to understand what is the demand for privacy, so at some point we have to specify some utility for privacy, so how should we do it?

DR. JOHNSON: Tough question. So I would say the following. I think your presumption is that you have to specify the classic economic utility function. And you can go ahead and do that. I'm not going to stop you. You just won't be describing the

process people are using when they're answering a question, can I gather personalized data.

So to use your example, it's known that zip code, plus birthday, you know, you know who I am. Do you think most consumers know that? So that's a case where there's a market failure and perhaps regulation is necessary. And to say people have a utility for privacy when they don't even know the basic facts about how the information is used seems -- I want to be polite here -- seems perhaps inaccurate. And not the basis of good analysis.

DR. GOLDFARB: Can I react to that?

DR. JOHNSON: Please.

DR. GOLDFARB: Okay. So, there is a difference between saying the utility of a full-information model and saying there's a fundamental thing called privacy that we care about. And we're mixing a little bit about privacy and security here, and we'll get to that in a second, but let's just talk about privacy and not worry about fraud and security, okay?

So there is a fundamental thing called privacy, perhaps, that people may or may not care about, and that is a utility construct. I don't know how to think about it in any other way. I'm an

economist; I accept that. But, you know, we're at the Bureau of Economics; I'm allowed to say that. And the -- that's different from saying consumers have full information.

And we have good models of thinking about information -- pretty good models of thinking about information. And, so, just because we don't know how to construct a full-information utility function doesn't mean we should throw out the idea that there's utility to privacy.

DR. JOHNSON: One last response. Years ago, there were proposals that you asked people what do you want to happen in certain situations. And, so, rather than every time I go to a website I have to sort of decide what boxes to check -- I think one was called EPIC, I forget what it stood for -- I was at FTC, maybe the 2000 conference.

So basically I say do you want other companies to know -- to sell your Social Security -- your birth date to other people, which, by the way, if you get the zip code, it's the same thing. But, you know, and you basically would make that decision once with information, and then that would be captured in a profile that would be carried to every website you

217

went.

Now in a world of assembled preferences, that seems like a much better way of intervening than, you know, assuming that I can regenerate that every time I visit one of the 30 websites, 50 websites I visit every day. So part of it is it does matter when you come into action what are the interventions. And, so, if you help people assemble functions -- utility functions in a way they won't regret, I think that's sort of one of the interventions.

AUDIENCE: I think the problem with that, I mean, that's why Facebook is so popular. It acts as just a gateway. And --

DR. STIVERS: If you can wait for the --

AUDIENCE: Oh, sorry. I'm still recovering, so this helps. The problem -- I mean, this is one of the reasons, you know, Facebook has become so popular for the sign-in because you don't need your credentials, right? You just use Facebook.

But then the problem is people don't know that they need to go through those arcane menus to uncheck and they're passing a lot more than their zip code and their birth date, right, and all their preferences. And, you know, I'm sure you do this, and when we talk about this in our digital and social

218

media classes, students are -- have no idea this information is being shared. But then after you tell them, they don't change it.

AUDIENCE: I was just thinking about what you were saying, Avi, when you talked about the Fair Credit Reporting Act and its beneficial purposes. And this is a little different than what we've just been talking about for the last minute or so. Could it be that one of the interesting differentiating aspects of that is the existence of some third-party non-individual-aligned entity where data resides?

You know, as sort of an electronic ombudsman or intermediary? I was keen here on your idea about innovation. If firms want to innovate services and products that actually people want, they do need to know more about people and what they want, but maybe individuals don't want to reveal that.

So if we could have third-party ombudsman like repositories of information about cohorts of people who are willing to be -- to put their information in, you know, that might create organizational structures where some data could flow. It's just a crazy little idea, but it would create an anonymized database in the same way that we benefitted from Nielsen data forever and ever and things like

219

that. Just my little idea.

AUDIENCE: That's actually happening in the credit card industry now. Third-party -- third parties are being created that are allowing all these banks and credit cards acting as the intermediary to -- and stores can share their information into this database to then -- it's all anonymized, but then you can pull it out, just exactly what you're doing.

AUDIENCE: So you were just saying, doing that, yeah.

DR. STIVERS: I think in both of these -- with the credit and credit cards, one of the issues that I think was hinted at -- or maybe even said explicitly by both Eric and Avi -- is this idea that accuracy is actually something that consumers care very much about.

So if you have my information, if you're going to be acting on my information in some sense -- and credit is one of these issues where you basically are going to be acting on that -- that's going to be, I think, potentially a way to, A, make consumers pay attention to, hey, what is this information going to have, but also to alleviate some of this, well, hey, I want to be really private. Well, but I also want you

220

to be really accurate in terms of how you address your decisions toward me.

AUDIENCE: So, people have been talking about privacy as if it were a light switch almost. Do I want companies to have my information or not? And something I'm curious as to why nobody has mentioned, either on the panel or in the audience, I might say, do I want companies to have my information, no; but I might be willing to sell it to them, depending on -- depending on the type of information.

Even in the extreme case of security, I probably wouldn't sell my Social Security number to anybody, but I might sell my medical records. And -- for a higher price than I'd sell my favorite color, but the issue of consumer willingness to charge, I just -- I want to open -- anyone want to comment on that?

DR. MAYZLIN: I mean, Alessandro Acquisti has done some experiments on the value of privacy, and I think often it looks -- I think in the lab people say that they care a lot, but when they -- you know, revealed preference says they don't care at all. I mean, when people, you know, put all these things up online, you know, and don't have very good privacy controls set up, they act as if they don't really care.

So I think there's -- right, there is a big disconnect between what -- you know, what -- like if how much you want to pay for your medical records, you know, is $100,000, but then you'd basically put it -- you know, you talk about every time you're sick; you talk about -- you know, so -- so I think that's the kind of weird thing about this field.

DR. GOLDFARB: So I would add a couple things. First, so we can think about a property right to the information, and that's where this would go. And you say the property right lies with people, and then they can sell it or not. If we take Garrett's results or his speculative -- you know, his ballpark numbers seriously, it's hard to think of the transaction costs of thinking through that market being sufficiently low that we can justify that any trade will happen.

Maybe, you know, there's lots of great technologists in the world, and maybe eventually we'll get there, but that's, I think, a first-order challenge. And there's a second challenge, which is because information is non-rival, it becomes, once again, hard to enforce that property right in a way like you try to enforce copyright and it's hard

enough. And there the incentives -- the commercial incentives are much, much higher.

DR. STIVERS: Eric, if you have a quick --

DR. JOHNSON: Very quick. So if this is an assembled value, the following should be true, and I bet you it is. I say you're going to buy the right to keep your information private versus you're going to sell the right. We know that from mugs that's two to one. For taboo tradeoffs, that's zero to infinity. I expect it's going to be closer to zero to infinity than it will be to one.

So I don't think that value exists, although this idea is great and Esther Dyson was talking about it in 2002, and just -- the market never has happened for reasons I think Avi's right.

DR. STIVERS: Well, I want to thank our --

AUDIENCE: In the case of medical records, for example, there is a black market that's -- in the case of medical records, I know, for example, that there is a black market on which any of our medical information could be bought and sold that was hacked from our insurer.

DR. STIVERS: Okay. Unfortunately, I do need to cut us off. I do want to thank our panelists for participating. Privacy and data security seem to

be this kind of very big gravity center, so we always -- or often tend to go there. There's a number of other topics that I would have loved to delve into that were brought up by our panelists, but we're out of time. So thank you very much.

(Applause)

DR. JIN: Thank you. We'll be back at 1:50.

SESSION THREE:
ALGORITHMIC BIAS? A STUDY OF THE DATA-BASED
DISCRIMINATION IN THE SERVING OF ADS IN SOCIAL MEDIA

DR. JIN: Hello, everyone. We're going to start.

Thank you for coming back. I know the room is freezing, and we're trying to correct that. Okay, just give us a little more time. Hopefully, we'll be able to get it right.

So our next session has three papers. The first one will be presented by Catherine Tucker from MIT on algorithm bias.

DR. TUCKER: Okay, so thank you very much. So, this is joint work with Anja Lambrecht, who is sitting there wrapped up in multiple cardigans. And we're incredibly excited to present this today. This is the first -- a very new paper, first time we're presenting. So I'm going to try and go quite fast -- it's a very simple paper -- so we can get lots of feedback.

Now, what is this paper? Basically, we're going to use data from a field test and then go on to delve into whether or not it's suggestive of algorithmic bias. Before I go any further, I should say that -- I mean, I'm a marketing professor and therefore very, very proud of having worked with many industry associations, and, all the tech firms, apart from Apple. Now because this is the FTC, I should also make clear that this research was not funded by anyone apart from the NSF.

Okay. So let's go onwards. So our research question is to basically delve into the why and to start to present -- I think present -- some evidence about why it is that an ad- serving algorithm might appear biased. Now, why are we doing this? Well, we're doing this, like you heard during the panel, my gosh, we saw the privacy debate there, and I was recently at FTC PrivacyCon, and let me tell you, marketing professors, we should all be there, it's a wonderful conference. I feel we've got a lot to say.

But one thing which really struck me about that conference was the extent to which -- although the privacy debate hasn't -- is not just focused on the question of whether companies should be allowed to amass data; it's also now concerned with the question of, well, what harms potentially could firms do if they do amass data. And one of the most highlighted harms that could happen is basically the potential for firms to use their algorithms and all their data to

potentially act in a discriminatory way against individuals.

And, indeed, at that Privacy Con conference, there were two papers which looked at ads, and they both suggested that perhaps ads which might be desirable were often less shown to women. And there was also a paper which looked potentially at ads being served less -- certain different ways depending on race.

Now, those papers were basically documenting a pattern. And what we aim to do is build on that literature and actually look at why. Why is it there's an ad-serving algorithm that might produce effects that make us feel uncomfortable? So what we do is we have data from a field test on an ad which promotes job opportunities in the STEM sector -- for those of you who are not familiar with that term, that's science, technology, engineering, math -- and this ad is going to be shown across 190 countries.

And the ad was set up as being gender-neutral; however, it ended up being shown to more men than women. And we might think, well, is that a desirable outcome? No. Especially given that the STEM sector is a set in particular which has struggled to attract women.

So we have that result. We've sort of got this headline effect, 20 percent less likely to be shown to women. The question is why, and that's why we think we're different in what we're doing in this paper in that we show it's not to do a click propensity; it's not the case that women just don't click on the ad and the algorithm is reacting to it. It's not the case that there was less opportunity to show the ad to women because they're on social media as much. And it's not the case that the algorithm had learned some kind of underlying sexism from the host country.

Is that what we show, that in some sense what we're seeing is very much unintentional bias in that young women are a valuable demographic for advertisers? As a result, it costs more to show ads to them. And, so, if you have an ad algorithm, which is just trying to minimize costs, then that can lead to a situation where the algorithm shows fewer ads to women.

So why does this matter? Well, it matters - - well, what we claim is that we're the first paper to really sort of look at the why of why we might see these adult-serving algorithms serve ads in what appears to be a biased way. And what we show, which I

think is quite intriguing, is it's not the case that we have an evil ad algorithm. Instead, we have an ad algorithm behaving in a way which might look biased on the face of it which is the result of a series of completely independent advertiser actions.

And one thing I just want to take, and I'm going to riff off Avi's talk earlier, is that, as you've seen, the way we've always thought about privacy in the legal debate at least and the legal conceptualization of privacy is so often focused on the individual. And as you saw the definitions of privacy have focused around an individual. And I think one thing this paper does is highlight the extent to which we should think of privacy online as often -- or the potential of privacy harms as often being the result of integrated decisions.

Now, why is this important? Why did we send this paper, even though it was new to the FTC? The reason we sent it is that we know that the FTC -- this is something they're worried about. This is an article from PC World where they talk about this as being a -- you know, we've got people in the FTC in the room who can say if this is right or not -- that this has actually been a big topic of concern, especially among the technologists at the FTC.

And why do we think it's going to be hopefully somewhat useful for people thinking about algorithmic bias at the FTC is that at least has been discussed various policy solutions for algorithmic bias; one of which I heard a lot about is a solution called algorithmic transparency. And this has been sort of a big slogan, I think, for a lot of tech advocacy groups. And the idea is, well, maybe we could stamp out bias if tech companies just made their algorithms public, put them on the internet and we could study them.

You know, if you study this ad-serving algorithm, it wouldn't help you predict that it was going to react in this way which led to apparent gender bias. And, so, our paper certainly suggests that algorithmic transparency is not going to be a complete solution.

Another solution which is often discussed is potentially algorithmic auditing. That is seeing what algorithms are up to and just measuring the outcomes. Again, I think our paper emphasizes there needs to be a little bit of nuance there in that you could just do the algorithmic audit, but unless you try and understand why apparent algorithmic bias happens, you might, I think, unconsciously, unintentionally think

there was more bias perhaps than there actually is in the real market case.

All right, so, that's the basic motivation. Now I'm going to tell you about the field test, and I actually want to tell you a little bit about the origin of the test because it's quite a wonderful story. So after the FTC Privacy Con conference, Anja and I were inspired to basically try and echo some of the results that we saw. And we sent a huge team of undergraduates, and we sent them to work basically on gathering data about ad serving.

And now we had two groups of undergraduates. One was sort of the MIT group; the other group from Wellesley. And the MIT group did a wonderful job basically building something very complicated to scrape data; and the Wellesley team did something completely different, which was really their own initiative, and they did a field test. And as yet, we haven't quite worked out what to do with all the wonderful MIT workings, but the Wellesley thing is like brilliant and, like, you know, so straightforward, simple.

And, you know, we learned a lot from what they did. So that's the origin, and I want to give full credit to this bunch of 18- and 19-year old passionate girls from Wellesley who really sort of are

the inspiration of this.

And what they did was they basically found a website that talked about careers in science, technology, engineering, maths and created an ad that linked up with it. And the ad was very simple. You know, it's not going to win any prizes for advertising. It just said there were STEM careers; find out about them. That's the ad.

And the ad's going to be the same. We're not going to do any fancy things to the ad. All that's going to vary, and I want to emphasize, I've got to make sure I call this a field test, not a field experiment, because we don't vary that much, but we're going to target it at 191 different countries, basically the entire world.

We're going to make sure that the ad was shown to at least 5,000 people in each country. And, now, one thing I should just highlight is that when we set it up, we worked very carefully to set it up to say that we're going to target both men and women. We didn't say men or women; we said all. And the aim was to sort of try and at least choose something which on the face of it was meant to be gender-neutral.

Now, the only thing, as I say, that is actually changing in all our settings is just the

location. We're going to have 190 different locations, different countries, and we did this because we wanted to sort of delve into the question of whether an algorithm can unintentionally pick up on the bias of the country concerned. That's been discussed a lot. We're not going to find any evidence of it, but that's why we did just so many countries.

Now, I just want to emphasize, you know, as you can tell, a wonderful team of Wellesley students that, you know, they didn't have much money. NSF gave them a little bit of money, but we didn't spend much money on this. Why is that? Well, a lot of these countries, it just doesn't cost that much to show ads to them, right? It's not America. We have a little bit of a tail in terms of how much we are paying to show these ads, and that tail is basically driven by English-speaking, rich countries.

So let's now get -- I'm going to just point this out. This paper, it doesn't need complex analysis. I'm going to show you everything I really have to say in a series of just tables. I'm going to show you two tables to get cunning, but it's going to be quite simple. These are basically -- this is overall what happened. And you can see these are the number of impressions, and you can see there's a

233

contrast between the number of times the ad was shown to men and women. And you can also see how many times the men and women clicked.

Now, I want you to notice three things from this. First of all, it is the case, this ad was definitely shown to fewer women than men. And that difference is particularly pronounced, I would say for sort of quite young women. I would also say that, you know, if you look at the click figures, they look quite similar right on the face of it. I could show you this table in a different way.

And this table, you know, I could go per country. You know what you should see there is just basically the same pattern. It looks a little bit different on the country level just because the small countries tended to have fewer older people in them. They were a lot of Caribbean countries. But basically those three patterns hold.

So from the broad data, what you should see is that men see more ad impressions than women, headline result, particularly among younger, if you look at younger ad cohorts. But the clicks appear similar.

Okay. So, I just said this paper doesn't need any complex analysis, but because we're

234

economists, it doesn't mean we didn't try and do a regression; we did. And, so, the question is going to be incredibly simple. And basically what we're going to be focusing on, just in a regression format, is how both the binary indicator for female affects how an ad was displayed and also the interactions between that binary indicator and age.

Now, I'm going to make you squint a little bit because this is going to be our first results table, and it's going to be very like what you saw in those tables, but at least you get sort of an idea of statistical significance. We are going to find profound effects as indeed women do see less ads than men. And it is particularly pronounced among women of the 25 to 34 age group potentially.

Now, the other thing I wanted to highlight is that if we just look -- you know, you can see ad impressions just means the number of times an ad was shown. In some ways you might be a little bit more interested in the number of unique users we reached. If we look at that, in fact, actually our results are, I would say, almost stronger. And the reason they're almost stronger was there's this odd thing that if you were a woman who happened to see an ad, you tend to see it more than a man.

235

Okay. So, those are the three big results. You know, they're very straightforward, just there in the raw data. We're more interested in the why. Now, when I have discussed this paper, you know, sort of excitedly with people, you know, when you're trying to tell them your research, but they never let you get to the end and tell you what your result is before you finish, what they always said is they've always said, aah, I know why that happens; it's because women don't click on the ad. Right? This is a universal sort of idea that women are bringing this on themselves because they're just uninterested in STEM careers.

Now, let me tell you, that is not what happened. You saw it in the raw data. Actually, it's not the case that the ad algorithm's just reacting to the fact that women dislike this ad. Instead, if women see this ad, they are far more likely to click on it. Now, this isn't really moderated by age, but we do see in general that women do click more on this ad. So it's not an explanation; it's not the case that simply the ad algorithm is optimizing based on click totals.

So another potential explanation is okay. Well, maybe there are just fewer women out there to show the ad to on social media. Let me tell you, we

236

have looked long and hard and we can tell you every single piece ever written on this says women spend more time on social media. We will get -- you know, we will, if you like, pay lots of money to comScore to confirm that and MicroData, but really, you know, it's sort of a known fact that women spend more time on social media. So we don't think it's sort of constrained-ish.

Now, what we're interested in and why we did the 190 -- that's sort of one of the original ideas was this idea that unintentionally algorithms can pick up the bias of their host countries. And the idea is that maybe they have a training set; they've learned that over time that for whatever reason women don't click on this ad and so, therefore, that bias is how they show ads in the future. That's the idea, and that's what we're really interested in.

But you know what, we didn't find any evidence of this going on. Instead, when we put in interactions for -- at least sort of World Health Organization data about female labor market participation, it didn't pick up anything. So whether or not women were more likely to participate in the labor force, whether or not women -- I've got the result, but primary education in general, whether

237

women were more educated, that doesn't affect our results. And, also, if we just take one of these indices, which the World Health Organization constructs for female equality, nothing changes.

So it's not the case that what we're picking up is some lingering bias in the algorithm. Instead, the result is going to be -- the actual reason we think this is happening is so prosaic and a lot more straightforward in some sense, which is that we think it's really to do with pricing in a world where advertisers are bidding on an individual eyeball.

Now, if you look at a lot more data, you wouldn't really see this simply because we sort of see the same price we're paying per click for women and men, but remember, we weren't actually bidding that much. And, so, there's still the potential that actually what we're picking up is that we just didn't bid enough to reach women.

And, so, to investigate this possibility, what we did was we got the same wonderful team of Wellesley girls to go out and actually collect lots of data about bidding for women and men on Facebook. And this is something Avi and I have used before, but basically all social media platforms -- advertising platforms basically give you data on what you should

238

bid. They give you some suggestions. And, you know, we've made arguments in the past, at least Avi and I have; we got the paper published, that this is -- you know, it tells you something, right, at least about what the ad algorithm wants if you look at suggestive bidding.

So we got this suggestive bidding data figure for each of our countries. And this is what we found when we got this bidding data, and it was quite interesting. If you just collect this bidding data on average, women cost five cents more per -- five cents more. That's interesting. What's also interesting is that we wondered if this was actually perhaps itself echoing something about the value of women, so we also looked to see whether this was the result of cultural prejudice. This is actually more pronounced in rich countries. Women cost more in rich countries.

And then we went and, you know, we also got this data, so basically women cost more. Women sort of in these mid-tier younger sort of age groups in general cost quite a bit more. And, you know, that's particularly the case -- you saw -- look at the maximum bid, and you might think of the maximum bid here as picking up, well, what you have to do if you really want to reach that demographic.

239

So we think maybe our result was actually driven by the fact that women are just more expensive to advertise to than men, and so if you tell an algorithm that you're neutral and the algorithm wants to save you money, it's going to inevitably end up showing the ad to fewer -- to fewer women than men. And that's what we think is going on.

Now, the next question, of course, is why. Well, why do women cost more money? You know, we sort of started off this paper just waving our hands and saying, well, that's always been so. Women get married, have babies, those things cost a lot of money, maybe that's why. But then we realized we actually had some data which can help us sort of think about this. And this is data -- we basically got this huge data set of ads from social media, and this time this is consumer items, so it's an entirely different data set again.

And we're going to see how women behave about basically purchasing a wide variety of consumer items, ranging from vases to sort of decorative art. And when we look at this, we see some intriguing things which suggest we don't just need to wave our hands about why advertisers might pay for women; we actually see on social media platforms that women do

240

seem actually likely to exhibit behavior which might make them more profitable.

And, so, what I want you to notice here is it's not the case in general that women or young women, in particular, are more likely to click on ads. Now, remember, on our ad, they loved the ad. They liked clicking on our STEM ad, but in general, at least if you show them a picture of a vase, they're not more likely to click, particularly. On the other hand if they do click, they're more likely to buy.

In other words, in a world where you're paying for a click, if women are more likely to convert, they're going to be more profitable. And that gives us some rationalization about why it is that advertisers in general may be willing to pay more to advertise to a woman. So, in other words, what we're picking up may be something completely rational in terms of bidding behavior by advertisers.

Okay, so, let's get to the implications. So, there's a lot of limitations, of course. This is a simple field test, right? Very simple, straightforward field test. And this, you know, brings a descriptive to the words descriptive paper, right? It's very, very descriptive, intentionally so. What we do is we presented some evidence which we

241

think -- we rule out some things, and then we present some evidence in support of what we think is going on.

Another limitation, you know, we've got an ambitious title about algorithmic bias, but we only look at gender. And another big sort of limitation is there are lots of big -- what I would call the non-economist questions, which we don't tackle, in that, you know, I use the word "bias," but is it really bias when we have a world where an algorithm is simply responding to a lot of competitive bidding behavior? Should we call that bias? Should we think of it as bias?

That strikes me as a wonderful ethical question for a law professor. You know, also, should we think of this as discrimination? Again, I know at the FTC we probably have a lot of lawyers in the room, right, that's the sort of questions we don't try and tackle.

So, punchline -- it's not quite a punchline because it's going to have some policy implications later, but basically what we have done is we have this cross-national field test. And this field test, it was for a STEM ad. We tend to think of STEM as a desirable thing to show at least in a gender-neutral way if not trying to just because we worry about women

242

in the field; however, it ended up not being served to women. Instead, it ended up being served to men by sort of a figure of 20 percent.

We show -- but we show -- what's interesting about the paper is not just that result. But what we try and do to show why this happens, and that we show it's not to do a click propensity; it's not to do with local prejudice or the algorithm picking out local prejudice; instead, it just simply seems to reflect the fact that perhaps very rationally other advertisers consider younger women -- younger female cohorts to be a particularly profitable segment, and as a result are willing to pay more for them. And as a consequence, a algorithm which tends -- is intending to make cost-effective decisions on the part of the advertisers might end up showing fewer ads to women.

So here we have a nice -- I think a good example of a case where we have apparent algorithmic bias, but it's just simply an unintentional consequence, what I'm going to call external behavior. So what are the implications for managers? You know, some marketing business school professors have to say what managers should do.

Well, first of all, you know, since -- we've actually -- Anja was talking to some people -- some

243

advertisers in London about this result, and they were like, wow, we never thought of that. As soon as they thought of it, she said, don't worry, you can solve it quite easily, just run two campaigns for men and women. And they were like, we never thought of that. And, so, it was actually -- you know, so as you soon as say it, the solution is quite obvious.

So, good news, managers, there's something you can definitely do. But having said that, I do want to raise the question. You know, we do this because in some sense this is easy to do. Gender is an easy thing to look at, but there may be other types of bias that we worry about, such as race or economic marginalization, where we may see the same unequal distributions, but they're a little more difficult to measure and a little bit more difficult to know what to do about; so we want to highlight that.

Now, for policy, again, you know, we think this is an interesting case study where at least if we just looked at the algorithm it would just look like it's profit maximizing, trying to be cost-effective, no -- nothing about gender at all in it. So I'm not sure if, in this case at least, algorithm transparency would be helpful.

The other thing we want to emphasize is if

244

we go down to a world where policy is focused on auditing algorithms, then I think our study does emphasize that often what might look like it's a discriminatory outcome can be actually the consequence of potentially completely external or exogenous behavior.

So, with that, I will say thank you, and I look forward so much to our discussant.

(Applause)

DR. JIN: Thank you, Catherine. We love those creative Wellesley girls. And our discussant is Kanishka Misra from UCSD.

DR. MISRA: Thank you very much for the organizers. Thank you for having me as a discussant. And thank you for sending the paper. It was really fun to read. It's very simple and it's also as a discussant something you appreciate.

I am Kanishka, and this has been verified as a vendor, were trying to pose as me earlier; my ID was checked on my way up, so I'm definitely Kanishka. All right, so what I'm going to do is go quickly over the paper and then sort of pass on some thoughts.

There's been a lot of discussion in the popular press where it sort of headlines like STEM is a huge problem. Huge -- lots of articles in all the popular press about the under-representation and the gender bias in STEM careers. And recently in the U.K., there was a -- there was a finding that when looking at college applications, if it was a female name versus a male name, people viewed the application differently. And they're piloting a program where they're removing gender from applications, especially for the STEM careers.

In this paper, they're going to talk about algorithmic bias, and algorithmic bias here is defined as a advertising campaign that's meant to be gender-neutral, but it unintentionally was not gender-neutral. There were, again, another very sort of popular press algorithmic bias that came out recently.

This was from The Seattle Times, where someone found that if you go on LinkedIn and write Stephanie Williams, and actually it's true for many women's names, they come and say, well, do you really mean Stephen Williams, right? And that's -- again, the reason why that's happening is because there are just more Stephen Williams in LinkedIn's data set, and that's causing sort of this apparent gender bias.

All right, quickly, what does this paper do? This is a field test. It's a very simple ad. That's it, right? So it's a very simple ad which has do you think about careers in STEM. It was targeted to 18-

to 65-year-olds, not by gender. So if any of you have ever tried something like Facebook advertising, you input where you want to show your ad, the demographics you want to go after, you can go after different interests. You can also go in to say, well, which -- do you want to do men, women, or nothing.

They said the minimum bid was 20 cents, and as Catherine alluded to, English-speaking, rich countries and Switzerland had about three times higher sort of bid values there to do. Beyond that, they collected some good data from the U.N. about sort of gender equality in different countries. From that, what they found is women represented less than 50 percent. The numbers I'm going to show in my slides are for total impressions, I think, where what Catherine had was for reach. The reason I have total impressions is because they had that data in the paper.

So the main point is that even though this is -- they wanted to be gender-neutral, less than 50 percent of women saw -- less than -- women represented less than 50 percent of the reached audience. They said this is not driven by interest, and the data to support that is women represented 50 percent of the people who actually clicked on the ad.

All right, I just want to make a very minor point here. All right, so, and this is more the econometrician in me than what I believe, so I think -- I don't believe this was driven by interest, but econometrically you can say, well, perhaps the women who saw the ad are the women who are really interested in STEM and perhaps every other woman who didn't see the ad was not interested in STEM.

To truly get about interest, you have to get about, well, would the people who did not see the ad have clicked, and there's no way to get that answer, right? But I think it's completely fair to say that there's a continuous distribution of interest and there is not this huge dichotomy of it, as a conversion (indiscernible).

What they find is if you break down this 44 percent by age group, you actually see enormous differences across different age groups. The particular age group where women tend to be under-represented in their data were the 25- to 55-year-olds. And the question they're after is why. So what is happening? What's causing this? And why, even though an ad is clearly gender-neutral and targeted gender-neutral, why is this happening.

Interesting, they find no differences by

sort of this median split on the U.N. measure for gender equality, education, labor market participation. One thing I would really like to see in the paper, you always have -- create an amount of data, right? You have 191 countries, a huge representation of the globe.

It would be great to see more of the raw data than just a regression with a gender split, just to see sort of is there variation across countries. And, well, I actually don't even know -- is there enough variation across countries, and can something else explain it, if not a median split.

In order to find out more about it, they collected a different data set. This is a data set which looked at just the average price. Again, I don't know the platform they used, but again, if you order Facebook, you click whatever demographic you want to go after, and then once you do that, Facebook has a suggested bid or a minimum/maximum suggested bid. They looked at something similar to the suggested bid, and they find that's higher for women and particularly higher in the range of 25 to 44, which is exactly where they're under-represented in their data.

I went to some websites which suggest how you should advertise in Facebook, and I got this quote, and look, if lots of people want to reach someone, prices go up, not shocking, economic free market. If not lots of people want -- so it's very consistent with what they're saying.

They go one step further and ask to say why. Why is this happening, that women in sort of this particular age group are targeted by so many different advertisers. And for this, they collected even a third data set from a U.S. retailer, and they find the reason is because in this women in this age group are more likely to click on an ad -- or, sorry, from a click, add something to their basket and make a purchase. And that's what people should ultimately care about.

All right. Some comments. Firstly, the main results are very convincing. They're very clear in their raw data. They're very clear in their regressions, and it's very, very convincing, right? And they have multiple reasons for it. The answer -- the paper is very well written. It's sort of -- it's great that they collected sort of multiple data sets to make their point. And there's lots of sort of face validity to it. This is -- it's again the pricing argument, suggestive price is exactly similar.

I found a white paper which had a similar result but a very different headline. The result that women -- that men were cheaper to reach on Facebook. Their big headline was "men are cheap," which I don't know how I feel about that. They also cite other papers, which other sort of white papers would suggest that costs per clicks are higher -- higher for women.

Just some other thoughts of other reasons why more advertisers might go about -- go after targeting women. When I talk to advertisers, the number one reason I'm always given is women are decision-makers. This is true in a bunch of popular press. There's an HBR article about it, and this is purely saying that women make more purchase decision than men; therefore, more advertisers try and target women than men.

One thing you can potentially look at, and this is sort of a question which I had in reading the paper, that is this finding unique to social media or is this a gender finding? Do you find this in all forms of advertising? If it's driven purely by this nature of women make more purchase decisions, you should find it everywhere; if it's driven by something inherent about women more likely to click, then perhaps there's something different about sort of

social media.

I will make one sort of side note about this. So I did -- so for TV advertising, there are some planning -- there's some sort of suggestive planning websites. Thumbnail is one which I had data from, but my data are about 12 years old. Twelve years ago, when Thumbnail suggested how much you should spend to get a woman's eyeball versus a man's eyeball is exactly the same. So perhaps it's not true, but it's worth looking at.

The second thing they look at said -- they say sort of this -- and when we think about this as data-based biases, the data-based biases in this paper are unique to gender, but actually if you look at that data, there's more than just gender in the data.

So I told you when I was presenting sort of about what they did, they actually tripled their bid prices for three countries -- or four countries. So let's take a word where you do not have different -- mirrored campaigns by country. You have one campaign where you run it for the entire world.

What does this mean? Well, if you don't -- if you didn't triple your bids for these four countries, these four countries would be under-represented, right? And that's probably not data-

based bias. Like, we're probably willing to accept that. But, yes, it is more expensive to reach people in these four countries.

Also, I looked at that data. So I looked at their data. The orange bars here are just the impressions that they had by different age groups. I looked at the world population, and that's the blue bar, and then I took the world population adjusted by Facebook penetration. And, again, I don't know Facebook by website. I'm just taking an example of Facebook.

And what you do find is, yes, there is -- their population doesn't fully represent sort of the Facebook population and the world population. And is that a bias, or is that something that we're sort of accepting because of free market prices and they aren't bidding sort of very high amounts in their particular origin. So it's sort of an interesting question of, well, what do we consider bias and what do we -- or what are we willing to accept.

In terms of sort of main takeways from the study and why I think what we can learn, for advertising firms, as Catherine suggested, yes, there are differences. If that's important to you, you should mirror your campaigns, just like they mirrored

253

campaigns for different countries, if you want to reach the same number of men and women, you should have a mirrored advertising campaign or you have a different campaign to men, different campaign to women, just because the advertising prices are going to be different.

This problem is actually a very similar problem to what people face in surveys, what people face in polling, that it's just harder to reach some populations than others.

The second question is one for the advertising platform. So I actually asked some of my friends who work in advertising platforms, why don't you sell a way to buy ads where I can say rather than going after this demographic, I want this balance of demographics. The answer is they do sell it, but it's part of their consulting services; it's not part of the free thing you get access to.

For policymakers, I think the one big takeaway is that if -- and I think this is important and interesting to look at -- if you look at raw data and something looks like bias, it's really important to dive a little bit deeper to understand what's causing it, and maybe it's not biased, right? Maybe it's just something else, something in the algorithm

254

is causing it. And that's worth sort of thinking about before claiming or suggesting bias.

There's an interesting question about privacy, like what should be -- what should be sort of allowed or what should not be allowed. And I think especially when you talk about sort of under-represented minorities, that is something that we need to take a little bit more seriously in saying what should you be (indiscernible).

All right, thank you very much.

(Applause)

DR. JIN: Thank you. We'll take a few questions. Catherine, do you want to come over?

AUDIENCE: A very interesting paper. One question I had was if the algorithm didn't take -- reach as the criteria for optimization rather take return on investment, taking the clicks and buying into account, and given the fact that when women click more often -- I mean click they also buy more often than men, would this bias kind of get back to -- you know, minimize this bias because of the different kind of criteria that you are using?

DR. TUCKER: Yeah, that is an interesting question because I think one thing you can say that got results is, well, do you really mind that women,

255

men, same number of clicks, right? Is that -- isn't that neutral? I think the way we've been thinking about it is just in terms of equality of opportunity to have that click.

Now, I think what we could do about this potentially is there are different -- you know, you can tell an advertising platform to optimize different things, and we could potentially look at that, too. I mean, that's another way of getting at this.

AUDIENCE: I had a question about the variation across countries. And I know that you tried to do some of the United Nations index to control for differences across countries. But in the price that serves, you have only four developed countries. Part of me is not that convinced that the cost of reaching women is that high in many underdeveloped countries because the purchase -- right?

DR. TUCKER: No, that's exactly right. This is one of the things that -- you know, we had, I think, going into it, this -- you could -- I shouldn't speak for Anja -- but going into it, we had this prejudice that somehow the price for women would be worse in -- would be lower in less developed countries because they were less prized. But that's not really what we see. Poor countries, men and women are equal.

256

And it's the richer countries where women are higher priced.

AUDIENCE: Right. And, so, then, it has to be very clear that like, you know, in the other countries you won't see this bias, right? And, so, somehow I wasn't sure that I saw that, but it's nice to highlight that you don't see these bias in the countries --

DR. TUCKER: Oh, that's a really nice idea. So, you're saying you don't see this in Rwanda, but on the other hand, if we look at Taiwan, where for example there's almost a dollar premium for women, that's where we see it. That's a really nice idea.

AUDIENCE: Thank you.

AUDIENCE: All right, just to clarify, how was the campaign optimized?

DR. TUCKER: So it was done -- I'm trying to remember. It was -- we had a manual bid, so we tried to take a bid of it. So we had a manual bid, and we told the social media platform we were trying to get clicks.

AUDIENCE: Okay, click on it.

DR. TUCKER: Yeah.

AUDIENCE: Okay.

DR. TUCKER: Well, thank you. And can I

257

just say so much thanks for our discussant because what he didn't actually say was that he spent this entire week finding out all the mistakes that we made in the first draft of the paper and telling us, so he's just been amazing. So I just want to sort of give him a big shout-out. Thank you so, so much.

(Applause)

258

THE VALUE OF INFORMATION IN MOBILE AD TARGETING

DR. JIN: Thank you. We'll switch the order of the next two papers because Hema has a plane to catch. So the next paper will be presented by Hema Yoganarasimhan -- hopefully I got the name right -- from the University of Washington about the value of information in mobile ad targeting.

DR. YOGANARASIMHAN: Thank you. No, I was asking around, but no one told me it was actually eventually switched, so that's good. Thank you.

Okay, great. Oh, I have to speak in the mic, okay.

Okay, so, first of all, thank you to the organizers, both FTC and Marketing Science, for not just organizing this conference but also taking this paper. So it's still in pretty early stages, so, you know, this is a good opportunity for me to -- okay, mic. Okay.

So I'm hoping to get a lot of good feedback. I'm not good at this. So I'm hoping to get a lot of good feedback which might be helpful to the paper going forward. And I should say this was joint work with my first-year Ph.D. student, Omid, who has really been amazing in the kind of work that he's been doing. Faces are looking shocked and annoyed, but with

259

economic students, there is no karma this business. So, you know, so that's great.

So what is this paper about? So the paper is about --

(Applause)

DR. YOGANARASIMHAN: Yeah. I'll be fine. Oh, okay, it was the mic, not me.

Okay, so am I supposed to hold this throughout?

(Laughter)

DR. YOGANARASIMHAN: I see, okay. So what's the paper about? So it's about the value of information in mobile ad targeting. So we're going to look at what kind of information helps with targeting, how do you effectively measure the value of this information, and a little bit look at what are some of the privacy implications of storing and sharing this information. So that's really the goal here.

So let me start by giving you a little bit of background about the smartphone industry. I'm sure all of you probably know these numbers. So for me it was a little bit surprising when I first saw that there are 2 billion smartphone users in the world. I didn't know there were 2 billion people, so this was interesting. There are actually 7 billion people, in

260

case you guys didn't know.

So and the average, 18 percent, about 2.8 are on the iPhone, so that's quite a bit of, you know, internet usage through mobile phones. And much of this usage is coming not through browsers, as you might expect, but it's through programs which are known as applications or apps. Okay, and just to give you some numbers, again, there are about 25 billion iOS apps which have been at least downloaded once and around 50 billion Android apps.

So as you can imagine, then, given that this app usage is really driving this industry so much, both development as well as monetization of these apps is of interest to many players in this industry. So there are three really main or broad monetization strategies out there for apps. So what are they? The first is the paid model. So if you want an app, you go pay $4, $5, whatever when you download it and you can use it.

The second is, you know, what's known as now the freemium model where you can download a free version of the app which is basic, and if you want some extra features or a premium version, you're going to pay some extra money.

And the third is what we're going to focus

261

on, which is in-app advertising, and probably the most popular way to monetize ads where what you are going to do is you can go download the app, and it's free and you can use it, but every time you're going to use it you're going to be shown some ads. And that's how the developer is monetizing that.

Okay. Okay, so, let's talk a little bit more about in-app advertising, because that's what we are going to be really looking at. So I'm sure all of you have seen in-app ads. In case you haven't, here is an example. That little diamond that you see in the bottom is really the in-app ad. It's quite small. If you click on it, it takes you to the advertiser's website.

And, you know, just again some numbers about mobile ad space, it's about $13 billion, and I don't have the exact -- I don't have the exact number on how much of this is in-app advertising, but quite a big chunk is. So who are some of the key players in the industry here? So the first is, of course, the publishers who are the people making these apps and hosting them and looking to monetize them. And these are the people who are going to host the ads eventually.

And there are also the advertisers who are

262

going to bid and place ads. And most important player probably who is not visible to consumers is really the ad network, which is this two-sided platform which is going to match publishers and advertisers.

So one common goal all of them have here is that to increase ad response rates if they could, keeping everything else constant. So I wouldn't say, you know, these other things change, so it's sort of always the case that even if prices go up you don't want necessarily ad response to go up, but everything else being held constant, each player has some interest in seeing ad response increased. And how do we do that? We, you know, in marketing obviously the answer is that we do that with targeting.

So what is targeting specifically in this context? So targeting is basically you have an impression, which is a user who is looking at an app at a given point in time, and you have a set of ads that you can show this user, and which ad do you actually show them. So that's really the question that they're grappling with.

And, you know, there's been a lot of research in marketing, especially, you know, in the TV world and in the offline world on how to do effective marketing, but in the online or in the mobile setting,

263

you know, you can't do really demographic-based targeting or segmentation. So the two main ways to do targeting really is behavioral, which is going to use data on what the user did in the past, so, you know, what kind of apps he or she looked at and what ads they clicked on or what ads they did not click on. So everything about their behavior from the past. So that would be behavior targeting.

There's also contextual targeting, which takes into account not so much the behavior of the user but the context in which the impression is happening. So what kind of app they are in and what time of the day are they using the app and so on. So that would be contextual targeting.

So those are the variables on which you could be targeting, but there's also another factor which affects how well you can target, and that's the data that you have. So how -- you know, so if you're going to be training these models, what size of the data and at what level of granularity do you have it; is it very fine-grained or is it going to be aggregated, and what is the length of the status, so do you have one month, one year, and, you know, do you really need so long?

And, finally, depending on who is actually

264

doing this targeting or who is doing the bidding, you also have to worry about whether you can combine data across certain sources. So these are some of the things about targeting from the industry's perspective.

From consumers' perspective, of course, targeting can potentially be good because you see relevant ads in that case, but it comes with a certain cost because targeting by definition means that the advertiser and the platform know something about the user and, you know, and that's what they're basing, you know, the ads that they're being shown on.

So this is specific to the mobile context. So in the mobile setting, you know, tracking of users is actually very persistent. It's even more persistent than other online settings. So, for example, I don't know how many of you have done this, so if you go to your mobile devices, there's something called Ad-ID, which you can reset, but then you reset it, everything that you're doing through mobile phone can be linked across all the apps and across all the -- even the browser, I think, depending on how it's set up.

So a few years back, the system was actually even worse. So there was no way you could reset this

Ad-ID, so anytime you -- if you did something on your phone and a few months later you did something else on a completely different ad, these two actions could be linked using the device ID. So then Apple introduced what was known as Ad-ID, which Android also mirrored. So things are a little bit better now.

But now there's this question about, you know, apps merging data, advertisers merging data and so on as to whether these should be allowed and, you know, to what extent should even -- you know, we should have Ad-ID, should we even get rid of Ad-ID and not give this access to consumers so to advertisers.

So this is some of the background from the consumers' perspective. So what are we going to be doing here given this background? So from a substantive perspective, the first question we're really going to be looking at is we want to understand how much does targeting really improve the effectiveness of mobile ads, in-app ads.

So we really want to know -- we want to measure the target, you know, consumer response rate of in-app targeting. And that's from just a, you know, understanding, you know, how much does targeting help, but then we want to look at what if targeting were actually making consumers more responsive, what

is driving that, right? You know, what is substantive. Is it contextual information? Is it behavioral information? Is it a combination of both, and what kind of information?

And we also want to look at what's the value of data really and what's the value of more or better data in this context, right? And that's really for the standard perspective. From a methodological perspective, we want to really understand what kind of models perform well if you want to measure the returns to advertising. We want to look at econometric models, some of the standard ones, and compare them to some of the machine-learning methods and see are they better at being able to predict those.

And, finally, once we have some results in that, we want to then go and make a few changes in the system and look at two kinds of -- two broad kind of questions. One is what if tied into privacy regulations? What if you got rid of Ad-ID and you told advertisers on the platform that there is no more Ad-ID, you know, use some other metric to track consumers if you could. And how much worse off would we be in our ability to target?

And the other thing we want to look at is something which is not happening on the platform now

where the platform is not sharing much data with advertisers in terms of what the user behavior is. If you could allow more and more data sharing between the platform and the advertisers and between advertisers themselves, how much better could they target? Okay, so these are some of the questions we want to understand.

Okay, so given that that's what we want to do, what is the most -- the challenges? The first, of course, is that we really -- we really need a model with very high predictive accuracy. And standard econometric models, which focus on causality, don't necessarily work very well in this case, right?

So because what -- what those kind -- what these models generally do is you have some kind of a model, however nonparametric you might make it, and then given the margin of consumer behavior, you try to devise some parametric estimates. And what the problem that you are getting worried about is things like endogeneity concerns and so on because you're trying to make counterfactual predictions.

But when you are looking at a prediction problem, so we are not really trying to understand why you were targeting the effect there, right? So we are to some extent, but we are not really saying -- you

know, our bigger goal is to measure the effectiveness of targeting. So we are trying to make the best possible off-sample prediction that we can, right?

So in that case, when you need high out-of-sample predictive accuracy, your search space, from a modeling perspective, it's not just over parameters given model, but it's over actually the models themself, right? So that's really a tricky problem. Then you're running into things like bias-variance, tradeoff and so on.

So what this translates to is like when you want prediction, you have really working with a very large number of attributes, and when you fix the function of form and try to estimate parameters, you're going to get mediocre results. You want to also look at, you know, input for the function and form and the parameters, and even when you have something as simple as 38 features allowing two-way interactions is going to blow this problem up and make it into 1,600 features, right? So in the computer science language, you would call this an NP hard problem. It's not part of linear time.

So, you know, so those are some of the things that, you know, the problems that we run into, and that's why you see we turned to some of the

269

mission-learning algorithms.

Okay, so, there's a lot of related literature here, and many of the people who have all done this are all in this room, so that's good. So -- but unfortunately because of time constraints, I'm not going to get into all of them. But it was -- but it was interesting for me to notice that so many of the people who have worked on this are here. And especially I think Avi and Catherine. I think we were exciting because it was like A, B, C, D. I'm like, wow, how many papers, like, you know, in the same year and by the same authors.

Okay, so, with that, let me move on and talk a little bit about the data itself. Okay, so, the data comes from actually the major in-app advertising platform, as well as the App Store in Iran. Again, this is, you know, because of my very enterprising Ph.D. student, usually when I ask for data people always just say no. But it looks like when he asks for data, people always say yes.

So this -- as you might know, in Iran, American, you know, companies are not allowed to operate, but it's a very high-tech country, which means they have a very wide range -- local IP system. So think of this as a pattern of Google Play in Iran.

270

So this is -- they're both selling apps, as well as selling ads, selling, you know, have this platform to sell ads and our data is from the ad platform.

So it's about -- you know, I think it's one of the top three IT companies in Iran, and it sells over 50 million ads daily in mobile apps, so that's quite a lot of ads. And they have about 25,000 apps and 250 ads, and this ad site is growing quite a bit.

Okay, so, let's talk about the data and sampling. So what we do is we focus on the top 50 ads and the top 50 apps, which is approximately about 80 percent of the impression. And because of the model that you see, what we need -- we use, we need to sample the data. We are not going to use all of it, but the sampling is going to be over first -- or with users on a three-day framework. So we are going to take two days for training and one day for testing, and that's over about -- about 27,000 users. And which translates to about 17.7 million impressions for training in these two days, and about 9 million impressions for testing.

But to actually do this training and testing, you need a history of information, right, which is the features that you're going to target these people on. And for that, we go back and look at

271

all the history from one month before, and that is about 135 million impressions that we work with. Okay, and that's what we use to generate the features.

So what does the data actually look like? So the data is -- if you looked at the raw data, it's basically going to be for each impression it's going to tell you this Ad-ID, which is the user-resettable, device-specific ID. So until the user resets it, we know that this is this person, okay? And every time they reset it, it's a completely new ID.

We also know what is the app in which the impression happened, what was the ad that was shown. And we also have interesting education, which is the IP address of the person or the phone, which was being used. And we know of the time at which the impression happened, as well as the click indicator.

Okay, so, now let me talk a little bit about the framework. So that's the data, and that's what -- you know, we talked about the data now let me tell you a little bit more about what we do. So before I talk about the model, I just wanted to define the problem formally. So the problem is one of prediction, which is to accurately predict the property that an impression I, by a user U, in app B for an ad A at a given point in time with some global

272

history H will lead to a click, right? So this is what you're trying to do.

So, then, the goal is to devise an algorithm that takes as input a set of preclassified data, right? So this is the data for which you know that this is -- these are the impressions and at least some led to clicks, some did not lead to clicks, right? And to generate an output probably which is as close as possible to the true click property as in the test data, which is a completely different data from the training data.

So if you want to write this algorithm -- sorry, then what do you need basically? You need, I think, three sets of input apart from the data. So one is you need an evaluation method. You need something to tell you how well you are doing, right? And you could come up with many different evaluation methods.

The second is you need a feature set, and this is the -- what in marketing -- in the standard parlance we often call attributes or explanatory variables, right? So it's a set of features. And, then, finally, you need a classifying algorithm. And because we have our training data where we know the outcomes, this is basically a supervised learning

273

algorithm, right? So that's what it is.

Okay, so, briefly about each of these. So the evaluation metric we use is quite standard. It's basically the -- we take the log loss to begin with, which is, you know, some people often also call it entropy, but comparing log loss even across different data sets can be tricky because the baseline measure of how many clicks in a given data there is could be very different.

So you want to normalize it by how much -- you know, if you had to make, like, just average prediction, right, out of 101 -- you got one click, which was one person, and how much better can you do with your model, right? So that's why we do something like relative information gained which, you know, normalizes and based on like a completely uninformed guess that you could make.

So that is the evaluation metric we use. And to generate the features, we use a framework for feature generation, and this is something that I based on one of the papers that I worked on in the past. When you generate features, you run into this problem of, like, exponentially like expanding number of features, so you have to keep track also instead of that, that's why I used this functional framework.

274

And this functional framework works where it takes this input -- each function basically takes three or four inputs, which it relates to something about the data -- the user, the ad, or the ad or the time in which the impression happens, right?

And we had these features based on the impression, based on clicks, based on click-through rate, based on the variability in all the variants and how many ads people have seen or how many apps they are using. So this is one way to actually generate the features, but once you've generated them, what's useful to do is to classify them as behavioral features or contextual features or potentially both, right?

So, behavioral features are features which are based simply on user behavior with pure behavior features are ones which are based on user behavior with absolutely no contextual information, like contextual features that I'm going to get contextual features which might not necessarily have behavioral information, and then there are, you know, features which can do both.

Okay, so, now let's talk briefly about the classifying algorithm. So you can -- you know, there are zillions of classifying algorithms out there.

275

Okay, wow, five minutes. I'll be quick.

So obviously there is OLS, there is logistic regressions, and the one we use is boosted trees, which I'll explain in the next slide is a boosted version of CART -- classification and regulation trees -- and I've used it before in an earlier paper, and it worked very well. I could beat a bunch of -- beat out a bunch of Kagglers in the prediction context. So I assumed that if I could pick out a bunch of like computer science Kagglers that you don't -- this has worked well, it turns out my hunch was right.

And we also have a chapter on using machine-learning methods in marketing where we explain this a lot more, in case you are interested.

So what is the brief one-slide overview of MART? So it takes classification and regulation trees, which are essentially just trying to classify the data in very, very simple way, multi-conventionally, and tries to boost them, which is like add more and more of them to reduce the prediction error as we add more.

The nice thing about MART is that it does automatic variable selection. You know, you're not working with the 1,600 variables. It does

276

automatic variable selection, and it can incorporate, like, lots and lots of variables in a nonlinear way. And it has been empirically shown to be the best classified in the world, so that's -- so that's one of the reasons we use it.

Okay, so, that's the framework we use. Let me talk briefly about the results. So what this table shows is on the rows it shows the different methods of classifying algorithms. And on the -- in the columns, it shows what are the features it takes as input, okay? So the top row is all from MART. And you can see that basically MART outperforms, you know, the baseline prediction, and the logit models and OLS models, by a very significant amount, so margins -- you know, completely like, you know, beats them, so that's -- so that's one thing.

The second thing is when you look at the features, then, so what you know is that behavioral targeting is much more than just even pure behavior targeting which is the first column where you throw out all the contextual information; it's still much better than, you know, this pure contextual targeting.

Of course, when you combine both, you are much better off, but what it's telling you is that user-specific information is more valuable than context-specific information. And within context-specific information, we find that app-specific features are much more valuable than ad-specific features, which means that what's -- ad -- you know, one interesting thing about this platform, which I didn't have time to talk about, is that all ads end up being shown in all apps. So the ad itself is not necessarily very informative. We also thought maybe because the ad is very small and there is not much information in the ad, maybe that's why the ad itself is not giving too much information.

But, you know, apps seem very informative. And overall model prediction is pretty good, so you see about a 15.2 percent improvement in predictive accuracy compared to like a baseline where you are just making an average guess.

Okay, so, now let's take this model and try to think about some of the questions that we had earlier. So first question we really had is what if you got rid of Ad-ID, which is always a discussion which is happening, right? Why do you want to track people using this special ID, in which case they will be forced to rely on IP addresses, right? So that's the first question.

The second question we wanted to ask is,

okay, now what if you actually weaken privacy regulations, which the platform has put in place, which is that as the platform, if I want to allow the sharing of data with advertisers at different levels of granularity, what would happen. And once you start allowing -- once you start sharing data with advertisers, they could share it among each other. Then what would happen? Okay, so, that's the second question.

Okay, so the first thing is the value of users identify as Ad-ID versus IP address. So we do notice that if you moved from Ad-ID to IP addresses, you are going to be worse off. And significantly worse off. Unlike ad IDs, IP addresses change automatically. These ad IDs, you have to go to reset; for IP addresses, that's not the case. You move from one network connectivity to another, that's going to naturally change.

They are also going to be masked people behind VPNs are all going to show up under the same IP address, which means that you're -- you know, pooling all these users together, which is bad. So we find that actually getting rid of Ad-ID would be bad from a targeting perspective. And it goes from about, like, you know 5 percent loss is what we find.

So, now, the next one, if platforms are allowed to share data with advertisers what would happen. We considered what would -- what's the best case scenario now because as we get data on what is the ad-specific CTR. So we look at what happens if they had access to better data, which is ad-app-specific CTR. They told you, okay, in this app, this is your click-through rate.

And what if we actually gave them all your individual-level data for all the ads that are shown to you, for your ad. And then if we gave them -- this is the scenario for which is I think the most interesting, when you give them data from their own ads and give them some kind of cookie kind of information, with just your like history, without the actual individual-level data, right? And the fifth one is what, of course, advertisers really want. They want access to all the data, right? So, there will be an outcome there, what happens in this case?

So, of course, what's interesting is that we find -- well, we find that at least privacy-preserving arrangements are the best in terms of targeting, we get very close to it by preserving ad user privacy, which is that the scenario in which we give -- as you can see, if you compare scenarios 4 and 5, they are

very similar. So very few actually hold back the individual-level data from advertisers across ads but give them their own advertising data and give them this feature set, which really does not tell them much about the user once they go out of the system. Actually, you do get reasonably close to the first best scenario. So if you're able to do -- show that you can, you know, maintain privacy at the same time, maybe, you know, get reasonably good targeting.

And we also look at which advertisers benefit. We find that large advertisers actually benefit the most from these, followed by smaller and medium advertisers. The ones who control for the size of the data, the variation in the data is what it has. So even if you're a small advertiser, if you see a lot of clicks in your data, that tells that your data is more informative. So it's not just, you know, if you're large then you benefit more from this; if you're small and then you have a lot of variation, you might actually benefit more from this.

Okay, so, finally, we asked this question of what if you allowed advertisers to share data now. So each advertiser has access to their own data and now if they could share data with each other what happens. And then we look at the sharing past. We take each --

you know, the top 50 advertisers and we pair them with each other and say now if you pool together our data, how much better could we do compared to just, you know, using our own data.

And, again, we find that larger -- so, here we find that larger advertisers gained less from sharing because their own data is reasonably informative. But -- and -- but we also find that when both advertisers are advertising in similar contexts, their sharing is much more valuable. But one of the things we persistently find is that incentives of the sharing pairs are not perfectly aligned. So one always benefits, you know, significantly more than the other, which means that even if you allow this kind of data sharing they might choose not to do it because that is not an incentive-compatible payment system out there.

Okay, so, I think I'm out of time. Actually, I can see you're nodding vigorously. So what does this -- I think we can all agree that targeting is an important decision in mobile advertising, and what we are really trying to look at, you know, it comes with significant privacy concerns, and we are trying to look at this and find some answers on how do you actually do targeting, how do

you measure the value of targeting, what helps, and we're also trying to look a little bit more at incentives.

Some of those, you know, unfortunately I could not present a little bit more on whether the platform wants to share data with advertisers. And, again, that also we find that it doesn't. So that's pretty much what I had to say. Thank you so much.

(Applause)

DR. JIN: Thank you. Our discussant is Sridhar Narayanan from Stanford.

DR. NARAYANAN: I'll just use this.

Okay, full thanks to the organizers for putting together a wonderful conference and specifically for asking me to be a discussant on this paper. Kanishka mentioned that it was, you know, fun to discuss the paper because of how clearly and easy it was to read. In this case, it was -- for me, you know, I'm not saying it wasn't clear. The additional thing for me was that it also made me, you know, sent me on this journey of reading lots of papers in the media that I wasn't -- you know, I knew a little bit about it, but I'm kind of vague on the details of it, so it was fun to do this. Okay, all right, and specifically referring to all the machine-learning

literature and so on.

So first I'll go through a quick overview of the paper, then kind of try and kind of give a very brief intuition for some of the algorithmic part, which Hema didn't talk about at a pretty high level and then go into some comments and some suggestions.

So here's the research questions that this paper is trying to tackle. So ad networks have a lot of information, historical information, and they can share this information at different levels, and they have either based on regulation or internal policies, they have -- you know, different networks have different levels of sharing of information with advertisers.

So the question is, you know, what is the value of this information, both to the network, to the advertisers, and specifically, you know, what this paper is trying to ask is specifically looking at the question of, you know, in terms of prediction of clicks by consumers, right? So what kind of information is valuable, and what kind of aggregation of that and so. And finally to whom, right? And, so, those are the kind of broad set of questions that this paper is trying to -- trying to answer, the standard questions.

The overall approach is going to be to build a prediction model for predicting clicks by consumers. And then use this historical information, in this case a month's information, to build a set of predictive variables. Hema didn't talk about this, but they actually did some work to try and figure out how much and look at the volume of information. Does adding more information actually help in any significant way. And the broad conclusion -- I'm jumping ahead a little bit here -- but is that, you know, that there's a lot of value in relatively limited information.

All right. And, so, the next step is that they're going to take -- compare different approaches, specifically a couple of, you know, common, go-to econometric approaches with a couple of -- with one basic machine-learning algorithm, MART. And I'll come back to that in a moment. And then compare different kind of information-sharing scenarios, you know, using the kind of model that they've used to predict clicks. So that's going to be the broad kind of overall approach.

Now, I'll kind of do a little bit of a detour talking about CART and MART specifically because, you know, partly because this was kind of -- I'd read about it, but it was good to get refreshed, and I thought I'd share just kind of broad intuitions that I gained from that.

So the problem that these kinds of algorithms are trying to solve, or at least one of the problems that they're trying to solve is that there's potentially a very large set of predictive variables. And we want to predict some outcomes from them. So if you try to kind of use some kind of linear or polynomial regressions, one of the kind of underlying assumptions is that there is a globally kind of valid relationship between these predictive variables and these outcomes, right?

And, you know, if you kind of tried to make it such that, you know, such that this assumption is relaxed, you have an incredibly large set of potential interactions, not just two-way, but three, four, five, 1,500-way interactions potentially that you have to kind of think about. And, so, it kind of becomes and an impossible problem to solve using those traditional approaches. Okay.

So in reality, in different sub-spaces of the data, you might have very different relationships that exist between the predictive variables and outcome variables. So what does CART do? Basically it recursively partitions the data space based on a

variable at a time. I'll walk though a simple example to show this. And what is the aim of this partitioning? Basically to kind of differentiate data such that the outcomes -- you know, when you do a partition, you want to kind of find a partition such that you have relatively homogenous set of outcomes within each partition but kind of different across partitions, okay?

It does this by looking forward without revisiting the prior partitions, and that's what is referred to as the greedy part of this algorithm. But, you know, the reason it's done is because this is actually -- it has been shown that this is an NP-complete problem; in other words, you cannot find a globally optimal solution, so you have to use some kind of approximations for this. So, basically the sequence of locally optimal solutions gets you, you know, hopefully close to that globally optimal solution.

So just kind of giving an example of how these relationships differ in different parts of this space is an example from, say, the presidential race, and I'm not taking any names of who aligns where, but if you think about, say, you know, one of the key variables that differentiates the two major party

candidates is in terms of preferences of workers is potentially this, and, you know, maybe -- I mean, I'm not basing this on any data. This is just pulled out of, you know -- pulled out of my hat.

But, basically, this is -- you know, if you look at, you know, one of the major discriminators might be race, so you might argue that, you know, if you're non-white then, you know, your preferences are very strong for one of the two candidates.

But within that, you know, the differentiation within the non-white category might be based on -- first on, you know, the biggest differentiator might be whether you live in a red state or a blue state, and then other factors might start matching.

On the other hand, if you look at those who are white, maybe which state you belong to is not the primary factor after race. The primary factor after race is education, right? And, so, if that's kind of the relationship, you know, capturing it through some kind of linear function or a polynomial function or even interactions will lead to -- you know, will quickly blow up into a very, very large set of inflections, so that's why, you know, these models are relevant.

Now, what does a MART specifically or more generally this class of decision trees called boosted decision trees do? Basically what problem that it's trying to solve is that classification trees have pretty high bias, right? Even though, you know, what are called shallow classification trees, which means that, you know, the number of steps that you are going down is actually small, you know, how you stop going -- before I go there -- how do you stop kind of going any further is by the setting of please set criterion of how many branches you're going to have or some rule based on kind of optimizing some function, some kind of cost function or something of that sort.

But what boosted decision trees do -- the additional kind of problem is that of overfitting, so you have a high bias, you have an overfitting problem. So what the boosted decision tree does is relatively straightforward, even though in implementation it's hard to do, is that it -- the basic inclusion is that averaging across multiple decision trees helps you out by kind of reducing the bias but also kind of reducing some of these overfitting problems.

And specifically MART, what it does, is that it kind of fuses kind of a data-based approach to kind of finding -- you know, going through the steps of going through multiple decision trees from one to the next. The basis for which one you go to next is based on kind of finding the path of where the descent and gradient of some kind of cost function is the highest.

All right. What are the main results of the paper, and this is pretty high level, but I think that it's a fascinating literature overall, very, very vast literature, something anybody who's interested in this, you know, can spend a lot of time going into it.

All right, what are the main results over here? If you look at the ad networks problem, first it wants to find -- you know, one problem might be finding a good way to even kind of -- or a good algorithm to kind of classify this information, and what the paper shows is that MART does better than the alternatives and, you know, that's -- you know, that's a pretty straightforward result, something you would expect.

The other kind of results are that while, you know, putting together all the information on, you know, who the user is, the app, that they saw the ad in the ad itself, other information obviously is very valuable within that kind of work. Hema referred to it as behavioral targeting variables; things that kind of identify the user and their exact behavior is

actually more valuable than the contextual behaviors. So that was kind of useful.

Now, the other kind of problem that this paper is trying to tackle and, to my mind the bigger substantive issues, that about information sharing. Okay, and so what they find is that there's the highest gain in prediction happens when advertisers are provided impression-level data on their own ads.

And, you know, if they're provided information across competitors, actually the gains go -- I mean the gains are lower because of the simple reason that when advertisers get information about their competitors it softens competition. So from the ad networks point of view, it's actually a worse-off idea to share all the information. So that makes -- it makes sense that ad networks, therefore, don't share information across competitors.

I'll kind of jump to some of the comments. First of all, it's an important problem from the ad networks' perspective, there's a live problem, you know, what kind of information they're to share. It's also public policy problem because of what privacy issues some of the marketing efficiencies that, you know, different sources of information causes.

Okay. The data are nice and rich. An

additional feature which Hema didn't talk about is that there is -- the specifics of the auction mechanism of this Iranian ad platform kind of induces its own set of randomness. I won't go into detail of that, but I'll come back to the consequence -- or one of the consequences of this in terms of interpretation of the results later.

The empirical work is very competent and the results are kind of interesting, even though they are kind of intuitive as a summary. So, you know, some of the suggestions, first of all, you know, in this paper, the first part kind of compares different algorithms, and I wondered whether it cannot be more comprehensive than this.

Now, one of the rationale given by the authors is that there is private empirical work establishing the superiority of MART. In specific, there is one paper that is referred to, but -- and there are more that I looked at as well. But all of those refer to very specific conditions and typically average across multiple metrics, right? So they're better but not necessarily for the kind of context that you're looking at.

So there's several other promising candidates, and I won't go into all of them but

there's a large literature in this, and we can talk offline about some of those. But I think that that can kind of -- one of the objectives of the paper is to kind of demonstrate an algorithm, and for that -- from that objective perspective, I think there's -- it can become a little bit more comprehensive.

The second point relates more to positioning and phrasing. I wondered whether this is really value of information, right, because all the focus is on clicks, but more clicks need not imply value. After the clicks, there is conversion, and so if you really take it to the ultimate goal of the advertisers or the networks, the data is, I'm guessing, not there. But I think to say anything beyond clicks, but I think it can be more carefully worded so that expectations are clearly set up about what you can and can't do.

The third point relates to this point about the randomization that is done in this auction mechanism. And basically what it does is typically in Google and other kind of platforms there's an auction mechanism where there's a ranking and, you know, a score which is generated, and the highest score gets to place their ads. In this case, what happens is that it's not -- and that's deterministic. In the case of the auction that the data is from Iran was for, it is probabilistic.

So if you have a high score, you have a higher probability of winning an auction, of getting your targeted ad. If you're a lower score, the probability is lower. That on the one hand is nice because it induces this kind of, you know, random variation, which is one of the critiques of the machine-learning algorithms, whether on causality and kind of alleviate some of that concern.

But on the other hand, one of the main results is that, you know, bigger advertisers kind of -- and smaller advertisers differ in terms of the value of information, but they also differ in the probability of their targeting rule actually being applied because let's imagine that the limits, somebody who has an incredibly high score, has a probability very close to one of their targeting mechanism working, and at extreme, somebody close to zero is entirely random.

So I wonder if, you know, there's differences across big and small that they're picking up is also not picking up these fundamental differences in how much I can interpret it as a causal versus noncausal effect. So I think kind of a little bit more care in interpreting these results would be

useful.

So overall, this is a nice paper. It brings in, you know, you know, it's sort of an expanding literature and using machine-learning tools, but I think it's a very relevant area, relevant policy question that it applies it to. The data are great, and, you know, applied in a careful way. I just think a little bit more comprehensive analysis on model comparison, a little bit more care in terms of interpreting the results will make this a really nice contribution.

Thank you.

(Applause)

DR. JIN: Thank you. We're about 20 minutes over our scheduled agenda, so we can pick up probably just a couple of quick questions.

DR. TUCKER: Okay, I just wanted to say, so this is such awesome work. I actually think it should be two papers, and I think this should be --

DR. YOGANARASIMHAN: I should get an A and a B.

DR. TUCKER: -- that's right. I would say the second paper should be about privacy, because you could just do so much with some of your simulations, especially you can answer questions along the lines

of, well, how long do we have to track someone, right? When should data die? You could also tell us so much about, well, why is it that an IP address is working so poorly. Is it to do with the fact that there's multiple people in the household. Some really -- well, anyway, listen, I'll tell you, I'll email you all this, but I think there's a wonderful privacy paper to be written sort of secondary, which can really answer lots of important policy points.

DR. YOGANARASIMHAN: Thank you. Those are great ideas. I hadn't even thought of any of them.

AUDIENCE: All right. It's a great paper, Hema, and to continue with the question raised by -- or suggestion raised by Catherine, I think if you talk about the behavior in a contextual targeting in in-app ads, if your data can have some user-level or app-level when they define context, it means the time and location, most of the literature, so the app may have some tracking users, longitudinal on that, to do some kind of location profiling.

DR. YOGANARASIMHAN: When I mean contextual, I'm talking about three kind of things. One is the app, where the impression is happening. The second is the ad that is being shown, you know, (indiscernible) presents the context, and the third is the time, and I

know you have done some other work in the context like, you know, where the impression happened and, you know, how crowded it is and so on. And so we don't have that kind of data.

AUDIENCE: Right. So the related question would be the targeting, do you know what's the targeting rule of the app? Maybe it's different from the --

DR. YOGANARASIMHAN: So they don't actually -- so at this point, what the platform is doing is they actually have this ad specific to it, so actually they're not talking any -- so they're doing a very -- almost you could say not that they're targeted except like an average of ads that they click to read. So they're not taking anything with the app or the time in which the click is happening.

And that's one of the reasons why they started working with us because they really wanted to look at, you know, if they did more targeting, how would things change; should they be doing more targeting. That's a bigger -- I mean, that's something I didn't get to, which is a really big question because it could soften competition if -- like a bunch of sharing data has shown.

DR. JIN: Any other questions? Okay, thank you so much.

DR. YOGANARASIMHAN: Thank you.

(Applause)

DIRECT-TO-CONSUMER ADVERTISING AND ONLINE SEARCH

DR. JIN: Thank you. The third paper will be presented by Matthew Chesnes from FTC.

DR. CHESNES: Okay, thanks for including this paper in the conference. This is joint work with Ginger looking at direct-to-consumer advertising and online search. The usual disclaimer applies here as well. These are our opinions and not those of any of the Commission.

So, first a bit of motivation. The U.S. is actually one of only two countries in the world that allow direct-to-consumer advertising of prescription drugs. And what I'm talking about are sort of the ubiquitous, you know, commercials you see on pretty much every commercial break anymore on TV as well as, you know, in magazine ads and newspapers and increasingly on the internet.

Prior to 1997, you were allowed to do this in the U.S., but you had to provide what was called a brief summary. And that brief summary was really not that brief, and the drug companies really didn't find it advantageous to advertise on television because you had to provide all the risks and benefits in a certain way. And it just -- it didn't really -- it was too cumbersome or too costly for them to do so.

So in '97, the FDA clarified its policy and essentially said that the drug companies could advertise on television and provide a major statement, just kind of like a one or two-line statement about the major risks and benefits, and then provide some other outlet for these companies to provide more detailed information, whether that be a 1-800 number, pamphlets in an office, you know, go to the library, or, you know, more and more to go online just to a -- to the drug companies or usually the drug itself would have its own website.

So following that change, prior to '96, DTCA was about $660 million, and as of 2010, it's now over $4 billion. It's actually about $4.5 billion today. It's leveled off. Some of the drugs -- some blockbuster drugs have gone off patent, which is why it's leveled off.

The FDA aims for a fair and balanced disclosure, you know, with this policy. So there's been a lot of research assessing whether that's being met, that goals' being met. There's also a lot of research looking at how DTCA affects patient visits, drug choice, patient compliance. A lot of authors are in the room that have worked on those papers.

So at the same time all of this happening,

of course, internet -- the use of the internet for health information has grown dramatically. There's a recent Pew research study which shows that, like, over 70 percent of the survey recipients use the internet for health information, to find health information, and 78 percent use search engines to start that search process. So, you know, this is all happening, growing kind of at the same time.

So drugs are kind of unique because the typical consumer may have limited information about drugs. They're complicated. There's a lot of different sides to them. So maybe getting them -- getting the information from multiple sources, including online, from their doctor, and from peers.

So we're trying to look at that link between advertising and search. Catherine has a paper that is closest to this area where we're trying to kind of determine, you know, how that -- how those advertising are affecting search, are consumers actually going online to find more information, and then is the FDA's, you know, policy, is it really -- is it really being -- is it really succeeding.

So there's this active debate on DTCA on the two sides of it. So, you know, DTCA, of course, is just, you know, informs consumers about the existence

301

of drugs, maybe prompts them to do research, talk to their doctor, and maybe eventually seek beneficial treatment, which maybe they weren't aware of.

But then, of course, the other side of it is that the advertising itself may be biased and emphasizing the benefits over the risks. You also have the fact that consumers -- you know, they don't directly choose their medicine. They have to go and get a prescription, so it may lead to overprescribing in the end.

So our paper tries to sort of shed light on both sides of this debate, and we're just going to look at -- it's going to be a fairly basic paper. We just want to kind of get an idea of does DTCA actually, you know, encourage consumers to search for this information and then dig deeper into that and say, well, what are they actually looking for, what are they actually finding. Are they looking on -- you know, on -- are they going to FDA.gov and getting that information? Are they going to the drug companies' websites? What are they actually -- what are they seeking?

And then we'll do something at the end where we'll look at heterogenous effects. So I'm going to look at different drug types and searcher types. So

302

drug types, you know, the main things I'll stress today, we'll look at what type of condition do the drugs treat, you know, chronic or acute conditions, the age of the drug, and then the insurance coverage. And I'll briefly talk about searcher demographics, but I'll probably run out of time before I get to that.

So what do our data look like? They're coming from the comScore search planner tool, three years of data, 373 prescription drugs. So we started with the Orange Book listing of all drugs, and then we essentially selected a sample of drugs that were either -- had some volume of search or some advertisements in our sample. And that really just cuts off the -- both of those limitations would cut off the long tail of drugs that just get no search or no advertising.

We cover the five large search engines. Monthly data on clicks, on searches, clicks separate for organic clicks and paid clicks or sponsored clicks. And what we observe is the overall number of clicks on a -- like on a given month or a given drug, and we also observe how that's broken out by entity, so comScore calls these things entities. Think of them as websites, but sometimes they're aggregated to a higher level. And then there's also comScore

303

provides searcher demographics for each of these terms.

So, importantly, a big caveat of the paper is we're only observing that one channel, right? We're only getting that search channel. We're not getting direct navigation or any other way that someone may land on a certain website, but we -- you know, we have some evidence that, you know, the search engine is the gateway to the internet. So hopefully that's not too strong of an assumption.

Advertising data comes from Kantar. We've got ad spending by -- on prescription drugs by month for that sample, and actually even longer than that sample. Overall spending and then also by media, so it's broken out very finely. We're going to look at essentially broadcast, print, and internet aggregated up, because that's -- those are going to be kind of the main categories that we're going to be focusing on.

And then some others, other sources of drug information from the Orange Book, National Drug Directory, and the MEPS to get information on prescription rates, insurance coverage, drug age, things like that.

So what does the search data look like? So

304

on the left-hand side of this pie chart is just the -- just looking at the organic clicks that we see. And you can see that general health websites like WebMD, places like that, they get the largest fraction, over 50 percent of the organic clicks, the brand name is large, there are some -- the producer site would be, you know, Pfizer.com; the brand site would be Lipitor.com, right? So that's the distinction between those two.

EDU tends to be health -- medical sites of universities. The dot-govs tend to be -- tend to be FDA and NIH. And then there's other sites which are just -- we classify as they look like nonhealth sites, that they're just things that don't go into these categories.

So what we do, and this is not a strict definition by any means, but we classify the pharmacies, the brands, and the producers as promotional sites. So, you know, they -- obviously, there's some information on promotional sites and there's promotional activity maybe on informational sites, but we think that the primary focus of these sites is promotion; and the primary focus of general health, dot-govs, and dot-edu sites is more informational; and leaving the other out.

305

Okay. And you can contrast that with the paid-click destinations, and of course, you see a lot more paid clicks on the brand sites and online pharmacies. General health is still fairly large on the paid side, so you're still getting WebMD and those sites are buying those sponsored links. And there are even some, you know, 0 percent for the dot-govs, but even the FDA does -- for certain drugs, they do actually appear in the sponsored links.

Okay, so, just -- it's a little bit misleading the way I drew these two circles. There's -- one of them is actually much bigger than the other, of course. So organic is about 91 percent of all clicks; paid are the rest. The informational clicks, 96 percent of them are from organic links, whereas the promotional clicks are -- a quarter of them are going to come from the paid links, which is -- you know, I think that that's intuitive.

So the reason I point these out is when I show the effects and regressions, you know, the marginal effects may show one thing, but when you put them in terms of actual clicks in the absolute amount, that's going to tell a totally different story.

And then, finally, if you just aggregate over organic and paid, 32 percent of those clicks are

306

going to be on informational websites and 16 percent are on promotional. Again, that's going to -- that's important when we look at the -- interpret the results.

So just a quick graph of advertising over time. We can see the television advertising in red really picks up after 1997 and then sort of levels off in sort of the late 2000s. And the internet is a growing fraction of DTCA. So television is about 60 percent; magazines are about 30 percent; and then the internet, I think currently, as of 2011, is about 6 percent of DTCA.

So just some drug attributes I'll just briefly mention. This is the typical drug in our sample, is about seven years old. Thirty-five percent are classified as chronic, and the threshold we use for chronic is more than five prescriptions per patient per year. And we've done some robustness on that -- on that number.

Insurance coverage, 76 percent, so about a quarter of these drug costs are coming out of pocket. That's coming from the MEPS data. And then on -- the average drug has about four prescriptions per patient per year. Okay, so just to give you an idea of what the sample of drugs looks like.

307

So I'm sorry, this table is pretty small, but I'll just -- I just want to highlight a couple things on this. This just gives you an idea of by drug type what sort of search activity are we seeing. So a couple quick things. For the type of drug, you see that there's slightly more search for acute drugs, you know, drugs that a patient's not, you know, taking all the time, maybe they're searching a little bit more. But the advertising actually is the opposite. There's actually about 50 percent more advertising on chronic drugs compared to -- compared to acute.

If you look at insurance coverage, you know, you see that drugs that have lower coverage, so this is just below the median coverage, tend to be searched and clicked more. So the story there might be that these consumers might be searching for an alternative source of supply if their insurance company is not covering -- not covering their prescription.

So these columns are the percent of clicks on paid, promotional, and informational sites. And I really don't see anything systematic across this by drug type. So there doesn't seem to be much going on here. And if you look at searcher demographics, age and income of the searcher, again, we didn't see much -- just in these descriptive tables -- about people

308

clicking more, less into paid versus organic or informational versus promotional. So that's why, you know, looking at these numbers alone, it was kind of hard to discern stories, so that's why we -- you know, the regression framework which we'll present now hopefully will tell a better story or a more convincing story.

So like I said, this is a very basic framework, so we're just going to regress log search on own drug DTCA and also DTCA and that drug's class in the previous month; and we'll control for drug fixed effects and month fixed effects. When I say search, we'll look at separately searches, clicks, and then break it out by organic and paid clicks. And then also in the second set of results we'll look at informational versus promotional websites.

Okay, yeah, so I'm controlling -- we've got fixed effects for month and drugs throughout all of these. So the first set of results, just look across the top line here. So these are all elasticity estimates, so we see about -- for a 10 percent increase in DTCA, between a .2 and .3 percent increase in search in clicks. There's a little bit of a bump for paid clicks, so, you know, .8 percent increase in paid clicks.

309

But, again, we know that because organic clicks are about, you know, 10 times as big as paid clicks, if you actually do the magnitudes, it turns out that the organic effect is even -- is about twice as big as the paid effect in absolute number of clicks, okay? So that's -- just be careful when you interpret those coefficients.

And we do see some spillovers from the class. So this is all DTCA in the class pulling out the drug itself, right? So this is just any kind of spillover from -- for drugs in the same class. And, so, you do see some effects, particularly on the clicks regressions.

And then we break out the different media, the DTCA across the different media. And, again, we're just going to look at broadcast, print, and internet. And here you see that the DTCA effect is really coming mostly from broadcast and internet. There is some positive effects on print ads, but most of it is coming from, you know, television ads and internet ads. And it's a little bit noisier when you look at the class effects. But, again, a stronger effect for the paid links relative to the organic clicks, and that difference is statistically significant from each other.

310

Okay, so, then we're going to break it up into promotional and organic, based on that classification that I showed you. So here what we see is we see stronger marginal effects, if you will, on the promotional sites compared to the informational. But, again, because informational is larger than promotional, the effects on total number of clicks is about the same, between promotional and informational. And, again, larger -- slightly larger effects on paid promotional and paid informational clicks.

AUDIENCE: In all of these regressions, there are direct fixed effects (off microphone)?

MR. CHESNES: Yes, yes. Direct fixed effects and year/month fixed effects. Yeah, so when we say query, that's really what we're talking about. It's all just drug queries.

Okay, so then in terms of heterogeneous effects, I'm just going to add one term to this regression, where we'll add interactions between own drug DTCA and some set of covariates, whether they're covariates, whether they're drug covariates or searcher covariates. So that's the gamma terms that are up here. So same fixed effects for month and for drug.

Okay, so, let me just show you the ones for

311

drug. If you just focus on the interaction lines, it appears that the effect of DTCA on clicks is lessened for older drugs, so stronger for younger drugs. It's lessened for chronic drugs, so, you know, so strengthened maybe for acute drugs. If you'll actually compare the coefficients on chronic to the overall log DTCA coefficient, they almost cancel each other out. So really all the effect is coming from acute drugs.

And then we get a positive effect on the low insurance indicator. So this is, again, just the binary below or above the median for insurance. So, again, this is consistent with that story that consumers are searching for maybe an alternative supply source to -- if their prescription is not covered by their insurance.

Okay, when we look at searcher effects, since I have a little bit of time, so these interactions are searcher age and searcher income, so this is provided by comScore. So here we see that older searchers. You know, some of the results are mixed. Older searchers are -- the effects of DTC are larger for promotional, less for informational, and then income goes the other way essentially. So it's lower income leads to less promotional and higher

312

informational.

So we were a little bit surprised by the income result. The age result, I think, is a little bit more intuitive if you think that, you know, older searchers may be more responsive to DTCA. They may be taking more medicine and things like that.

So just to summarize, then, so this is -- all the results are really right on this slide. So DTCA is associated with more frequent searches and subsequent clicks for both the advertised drug and we see some spillovers. And the effect is larger for paid relative to organic, broadcast and internet relative to print and promotional relative to informational. But, again, if you do it in absolute number of clicks, then the effects are much -- are much more similar between these different categories.

And then the heterogeneous effects show that the effects are stronger for younger drugs, drugs that treat acute conditions. So, you know, maybe this has something to do with chronic drugs and older drugs being more -- you know, these are drugs that are more familiar to searchers. Maybe they don't -- they're not as responsive to DTCA because of that. And then we get these results on stronger for low insurance and older populations and higher incomes.

313

So overall, you know, we don't do a full-blown welfare and welfare analysis because we're only observing the clicks. We don't see what -- you know, what goes on after. We don't do the conversion part of it either. We don't observe that. But we think that these results are at least somewhat supportive of the FDA's original contention when they came up with these guidelines for certain drugs and for certain sub-populations.

But it's not all good news. We still see some -- we see some clicks that are going towards these, you know, either paid or promotional websites. So it's a little bit mixed, but I think it's in general supportive of their intention. So thank you very much.

(Applause)

DR. JIN: Our discussant is Jura Liaukonyte from Cornell University.

DR. LIAUKONYTE: Hello, everybody. My name is Jura Liaukonyte, and I'm from Dyson School at Cornell University. First of all, let me start my talk by thanking the organizers for this wonderful conference and for inviting me to discuss this very interesting paper.

So let me start by first summarizing what

314

the paper intended to accomplish and what are the main results and why do we care about that. So the main question that the paper attempts to ask is whether exposure to DTCA advertising drives consumers to search online. And then sort of derivative, second-order questions are -- is what kind of information are consumers thinking and whether that varies by drug type and demographics.

Why do we care? I think this paper sets up the discussion really well by highlighting sort of the two sides of the DTCA debate. So one side is claiming that DTCA is bad, essentially that there are no incentives for the advertisers to highlight the risk and it tends to overemphasize the benefits and mislead the consumers.

On the other side of the debate are the people who are arguing that DTCA is actually good because information is always good. This type of advertising provides information about the existence of the drug; and then consumers can self-diagnose, match their own symptoms with the symptoms that are highlighted in the ad and then seek treatment.

So, if DTCA is biased, then having people to seek further information online is actually good. So from the policy perspective, this is a very important

315

question to be asking and very important question to be seeking answers to.

Another thing that wasn't pointed out but I just looked it up, it was very recently -- less than a year ago -- AMA came out with a very strict statement encouraging to ban DTCA. So this just sort of highlights this ongoing discussion and the ongoing policy relevance of the DTCA advertising that exists.

So just to summarize, the paper finds evidence that indeed DTCA is associated with internet search, and some people go to informational websites; some people go to the promotional websites. One thing that I did notice is that authors are very cautious throughout the paper, at least the way I read the paper, not to explicitly label anything as causal, but implication is there. So the reader is sort of left to wonder whether the results represent marginal causal advertising-induced search lift.

So I think this is really a low-hanging ball -- a low-hanging fruit is to sort of strengthen the discussion and to focus on the causality. So in what follows, I will try to be helpful in giving some suggestions for how to set up this discussion and maybe how to try some alternative specifications to strengthen the causal argument.

316

So I imagine that the -- I imagine the authors might face the typical endogeneity taliban, as we call it, during the review process. So let me try to set up the -- let me try to set up the standard advertising endogeneity concern that arises in the advertising literature. So, essentially, intuitively, we have the situations where brands may plan advertising timing with partial knowledge of the unobserved category or time effects, essentially something that is really important for the consumer behavior that brands observe but the researchers, econometricians do not observe.

So is it something that we should be worrying in this case? So let me sort of give you an example. I do not have the advertising data that you guys have, but I did have a free source of U.S. Google Trends. So we would be worried about the endogeneity. For example, here, I am graphing Chantix, which is a smoking cessation -- prescription smoking cessation drug, and I'm also graphing quit smoking search term on the Google Trends. And you can see that they're rather correlated. It seems like quit somking is one of the new year's resolutions, right? All peaks correspond to January.

The part that I highlighted with a rectangle is the part that I know the advertising expenditure was the highest for the Chantix drugs. How do I know it? Because that's the only one that was mentioned in the paper for that month as like an outlier. It was one of the maximum spends in your data set.

So if you're really regressing the search on advertising, you might be picking up just the fact that there is -- there is an interest in the market and the advertisers know that and so on. So should we be worried about that? Fortunately, this is actually something that is observed, right? So we could just include -- I'm sorry -- so we should just include as many fixed effects as possible.

So my understanding from reading the paper is that the authors stack everything on the left -- stack all of the searches on the left-hand side and then include one -- one vector of essentially month fixed effects. What I was wondering if you do have enough degrees of freedom to include drug-specific time fixed effects or direct category-specific time fixed effects, so essentially fixed effects for category and month interactions.

Why is it important? If you look at -- again, I'm looking at the Google Trends, and here I'm plotting searches for quit smoking and hypertension.

And you can see that they're kind of -- almost perfectly negatively correlated. So what your month fixed effects are picking is just an average of that. So if you could include the time-specific sort of drug-specific time fixed effects, I think that would absorb all of that endogeneity.

By the way, I have no idea why people do not -- are not searching about hypertension on -- during January. That's a very interesting empirical observation.

Another thing that -- so, again, let's try to put in as many controls as possible. Another idea that I had maybe -- and I don't know the extent of your data -- maybe market fixed effects would be possible to include. Presumably, the data is available, but I don't know how important it is in your setup.

So remaining endogeneity. So once we have control for all of these fixed effects, is there another endogeneity remaining in the (indiscernible) that we should be worried about? So as an economist, I've been trained to think that advertisers are actually placing their ads optimally, trying to maximize the profits. But now having really done a lot of thinking about the advertising endogeneity, I'm

really skeptical to -- I'm really skeptical that actually advertisers do do things optimally.

And I have worked with some companies, and I do know that they do not know what they're doing sometimes when it comes to advertising optimality. And we also have this paper in the QJE that is telling us that really the information is not there for people -- for advertisers to have the information of what is -- when the ads are optimal and not. The signal is just too weak.

And on top of that, there is -- there are severe contractual and institutional challenges that complicate the seamless optimization. So, really, I'm a believer that what you are picking up is actually causal, and I think you can develop that argument that it is causal.

So another couple of things. I think the paper would be much stronger if it had a little bit more of model free evidence. So one of the things that I was thinking is whether if you could employ diff-in-diff approach in any way by juxtaposing, let's say, search in the United States where you have -- where you have DTCA versus search for the same brand of drug in Canada, where you do not have the DTCA and whether sort of those deltas are informative for your

causal inferences.

Another idea whether you could look into juxtaposing branded searches versus the generic searches. And, again, just sort of anecdotally it looks like there is a variation that could be picked up that might support this. And, then, I have a couple -- I will just summarize it really quickly.

I think it would be also nice to talk a little bit about the microfoundations of causality. So you could develop the argument more carefully that shows these microfoundations. So we actually know, and it has been convincingly shown in multiple papers, that especially TV advertising, which the biggest effect that you're picking up is the broadcast advertising, actually causes almost immediate searches.

And we have several papers that show that, and you can see I have included one very telling graph that just sort of shows these huge spikes in searches right after the ads have been aired. So we know that this is causally happening, but because your data is so aggregated in the month level, you could sort of develop that argument to really convince the reader that it is causal.

And the last thing I'm going to say, I am a little bit interested in the advertising content and how that affects consumer outcomes. And one of the things that we know about the ad content is that informative ads tend to be not that interesting and tend to lead to lower overall searches, but informative ads lead to higher overall searches online for people who are interested in the advertised products, so for people who are in the market for that advertised product.

And one thing that I just did, I looked at my data which sort of has this variable for advertising mood, and it just seems very striking that the prescription ads are labeled -- the prescription ads are about 10 times more likely to be labeled as informative. So here's yet another mechanism for you to sort of have this causality story unravel.

So, overall, I think you could -- I really like the paper, but I would encourage you to sort of strengthen the causality story because I think causality is there, but it's just -- I haven't read -- I haven't really found the word "causes" in your paper.

So the encouragement is also to perhaps add a case study where you are looking at the more granular data to show the causation mechanism, and I

know that Google Trends is now realtime, minute by minute; and I know that Kantar data is available at the second level, and both have market-specific variation.

All right. Thank you very much.

(Applause)

DR. JIN: Thank you so much. I really appreciate the suggestions for endogeneity, which is a problem that Matthew and I struggle with a lot, and I'm hoping our future referees are in this room so that we -- we'll see the argument on our endogeneity problem is not such a problem in our paper. So, with that, probably just for a few questions?

DR. TUCKER: Yeah, I was just going to say with all respect to your discussant advisor, I just wonder if the informative versus persuasive distinction here is masking some really interesting stuff in that in particular what strikes me about your period is it was a period of an explosion of social media, user-generated content, all of these things. I would -- you know, we've seen various hypotheses in the literature, the advertising interacts with social media in a very different way. You have the data to look at it. It might be wonderful.

DR. CHESNES: If I could thank my -- thank

our discussant, too, for excellent comments. Cheers.

(Applause)

DR. JIN: Okay, well, we'll break now for -- let's see -- ten minutes, and then we'll come back at 4:05. Thank you.

(Recess.)

SESSION FOUR:

MIGHT I INTEREST YOU IN AN EXTENDED WARRANTY

DR. JIN: Thank you. Thank you for staying here for the whole day. We have the last two papers. The first one is going to be presented by Sriram Venkataraman from the University of North Carolina-Chapel Hill, who is going to talk about extended warranties. Thank you.

DR. VENKATARAMAN: First and foremost, thank you to the committee for the opportunity to present. And thank you to Ginger and her team here for being such great hosts. And thanks in advance to the discussant, Matt, for his comments. I realize it's Friday, and I tend to have this reaction when I take the podium, I clear the room, but I'm attributing it this time to a treatment effect, which is a Friday treatment as opposed to my presence here.

So, what am I going to be talking to you about? First and foremost, this is work with a doctoral student of mine at UNC. The research questions that we're going to be exploring, I'll formalize it in a few slides, but I'm going to be looking at certain substandard questions around extended warranties, and the empirical context for this particular study is going to be the U.S. auto

325

industry.

I've been told, and there's a well known saying that is imitation is the best form of flattery, and I'm going to embellish it a little bit and say plagiarism is an even better form of flattery. So no better way to describe what I mean by extended warranties for some of us who are less familiar than to cut and paste directly from the FTC's website. I'm assuming that if it's on FTC's website it's kind of pertinent -- a topic that's pertinent and dear to many of the folks here in this room.

And I'd like to draw your attention to a couple of components of the blurb that you see up on the screen. First of which is I'd like to draw a contrast between what I mean by extended warranties versus what I mean by traditional warranties. So traditional warranty is often referred to as manufacturer-backed warranties or factory-installed warranties. These are warranties that come installed with your car, and you don't have to pay additional monies for it.

Extended warranties, on the other hand, or extended service contracts in my particular setting, these are again insurance products that you buy, and these are optional. And you buy it at an extra cost,

326

and I'm going to show you in a little bit the premiums that on average people pay for these products.

Much like traditional warranties, extended warranties are also insurance products. The key distinction between traditional insurance product and an extended warranty product is going to be that there is going to be some overlap in what's covered. There's going to be some non-overlap in what's covered, the specifics of which we're going to be exploiting for the empirics that will follow.

Last but not the least, I'm not necessarily going to bias our opinion or, you know, expectations on what you're likely to see in today's presentation, which is the blurb says it might not necessarily be worth the price. I'm not going to be studying the question about why people buy extended warranties. I'll still speak to that in some handwaving way, but I'll tell you why that particular question will naturally fall out of the research that I'm undertaking today or showcasing today.

The empirical context, as I mentioned, is the auto industry. And why the auto industry? Well, I think given the composition of whoever is left in the room right now, I think it suffice to say you don't need any convincing that the auto industry is

327

important to the U.S. economy. It's the -- one of -- as far as an industry goes, it's a huge contributor to the national GDP, employs tons and tons of people. And for better, for worse, I've been drawn to this particular industry for a couple of years. I've been fortunate to get some papers through, not always, but we try, right? We continue trying.

The reason I studied the auto industry in this particular context is because it has a lot of similarity with the blurb that we just saw in the previous screen, okay? So to kind of draw out what I mean by that, so my far right, your far left, is the set of -- or the menu of manufacturer-backed warranties that you get with your product. So every new vehicle comes with two types of manufacturer-backed warranties -- bumper-to-bumper warranty and powertrain warranty, okay?

So bumper-to-bumper, on average, what it does is it covers all parts associated with the vehicle, hence the name bumper-to-bumper, apart from the parts that are responsible for or susceptible to natural wear and tear. Okay? On average, it covers the vehicle up to 36,000 miles or three years, whichever comes first. Once the bumper-to-bumper warranty expires, the powertrain warranty kicks in.

328

As the name suggests, powertrain warranty is responsible for all parts that are responsible for moving the vehicle. Okay.

On average, it's 72,000 miles or five years, whatever -- whichever comes first. So these are things that come directly with the product. If you want to buy supplemental insurance, i.e., extended warranties, they come in a menu of -- you have a menu of offerings to choose. I'm going to kind of group all of them as basically forming two flavors of extended warranties, one of which is regular warranties and the other one being comprehensive warranties.

The regular warranty is one -- and both warranties, for the most part, what they do is they extend your bumper-to-bumper warranty past the expiry of the manufacturer's expiry period. So if you buy the regular extended warranty, it takes you from 36,000 miles to 72,000 miles, three years to seven years. If you buy a comprehensive, it takes you up to 100,000 miles and seven years, whichever comes first.

Okay. Again, going back to the blurb that we had on the previous screen, it's basically extending your bumper-to-bumper warranty. So it overlaps in terms of what products are covered with the manufacturer-backed warranties as well. Okay.

329

And that's going to be critical for the exercise that will follow. Okay.

So why do we care about extended warranties in the auto setting? So here are the numbers that I wasn't privy to until I started researching this topic. So in 2014 alone, $14 billion was spent on purchasing extended warranties in the auto space. If I took a poll of people, and going back to the panel that we had at lunch, I'm told that the best way to frame a question is to ask people if you don't want to bias a question, then ask them would you refer this particular product or service to your friend, assuming that you are a better citizen if you're responding in support of a friend.

I'm sure if I posed that poll here in this particular room, most of you would say no chance in hell should anyone be buying extended warranties. Yet if you look up on the screen, 40 percent of the people purchase extended warranties. And when I say 40 percent, I mean in the context of the auto industry alone.

So naturally this is a question that's going to be pertinent to policymakers, and as marketing managers, this is a huge business opportunity for us. So I'm hoping these numbers alone should suffice as

330

evidence for some interest in this particular kind of research.

What's going to be very important for us for the empirics is going to be 86 percent of all sales of extended warranties happen at the point of purchase of this particular vehicle. Okay? That's the consumer side story.

So -- Tim?

AUDIENCE: Do you lump in the used cars warranties?

DR. VENKATARAMAN: Actually, this is -- to preview what lies ahead, all this is going to be used cars. The entire exercise is going to be used cars. There's a specific reason why I do that. Okay?

When it comes to the perks for the firm, 20 percent of margins or profits realized for auto dealers are through selling extended warranties. So just to put these numbers in perspective, the average profits that a dealership realizes through sale of a car, which has been the bulk of the interest of academic research, at least on the academic side, the average retail margins are around 2 to 3 percent. So we are looking at several-fold here.

When it comes to underwriters, 50 percent of the people who buy extended warranties never, ever

331

file a single claim. Amongst those who do file claims, the premiums do not necessarily -- or the savings are not necessarily commensurate the premiums that they're paying. So naturally from a policy standpoint, these statistics should warrant the question, why are people buying extended warranties in the first place.

I'm not here to claim that it's a bad investment. I might at some point, but not today, right? But think of this as insurance products, right? When we spend monies on our health care and we buy and invest in premiums, we have no expectations that at the end of each year we're going to be recouping the cost of our premiums. It's basically a peace of mind investment that we hope that it insures us against large cost shocks, unanticipated cost shocks in the future. So I'm going to take exactly the same position even in today's presentation.

So given the numbers that you see up on the screen, no surprise that auto dealers and underwriters are aggressively marketing extended warranties to us. So I suspect many of us in this room have been recipients of conversations at the end of closing a deal at the dealership or have received a place or something like this where they're trying to induce you

332

into purchasing extended warranties. Okay?

So to formalize the research questions, I'm going to be answering the following two questions. One is when the auto buyers of extended warranty -- auto buyers purchase extended warranties, and when I say when are they more likely to purchase it before the manufacturer warranty expires or more likely to purchase it after the manufacturer warranty expires, okay?

Why is this question important? Well, it could possibly inform or provide us some kind of inclusion or understanding of what underlying mechanisms might justify these choices or rationalize these choices. Once we have a good handle on possible mechanisms that drive these choices, as a policymaker, I might be interested in kind of using that as an input to assessing whether there's a need for a policy intervention, and if there were a need for a policy intervention, what kind of policy intervention might I design, and when might I actually introduce this policy, time the policy intervention.

From a dealer standpoint as marketeers, clearly this is going to be directly relevant for targeting marketing because as a dealership I can figure out should I be aggressively targeting extended warranties to consumers before the warranty expires -- manufacturer warranty expires or after. How soon before and how soon after?

So the empirical setting, going back to the question that Tim asked, the empirical setting that we are going to be taking to the exercise is going to be the used vehicle market. So why this particular choice of data? Well, go back to our question. I'm interested in studying whether people are more likely to buy before or after the expiry of the manufacturer warranty. So if I look at the new vehicle market, the entire manufacturer warranty is intact, so there is no variation that I can exploit. So I am left with no other choice, and naturally I'm going to be using the used vehicle market.

It so happens from a substandard standpoint as well, the used vehicle market actually forms the lion's share of all sales, at least in the U.S. So depending on which resource you trust, anywhere from 55 percent to 79 percent of all auto sales in the U.S. happen through the used vehicle channel.

For the purpose of this -- of this particular exercise, the used vehicle market offers us nice, rich variation -- natural variation that we're going to be exploiting for identification. And

specifically what I mean by that is there are some used vehicles that are almost in pristine shape that have all -- almost all of the manufacturer's warranties in place. Some are really, really old and have nothing. And then you have everyone in between. And that variation is something that we're going to be exploiting.

However, with used vehicles, unlike new vehicles, it also introduces a set of econometric challenges for us, first of which is no two used vehicles are alike. Right? So we have to figure out a condition -- try to control as much as possible the role of unobservables.

Second, there could be strategic sorting on the part of buyers. And what I mean by that is the composition of people who buy younger vehicles could be very different than the composition of people who buy older vehicles. And I need to find a way to tackle that as well.

Going back to the numbers that I outlined a few slides ago, if dealerships and underwriters are making tons of money on extended warranties, perhaps it's possible that the dealerships could be offering more attractive terms on the vehicle to kind of get you -- or get to win your business on the extended

warranty side. So I need to figure out a way to also control for those kind of possibilities in the empirics that follow.

So given the question, given the empirics, and given the threats to identification, the empirical strategy that I'm going to be taking to my data is going to be the sharp regression discontinuity design. So I feel like this design is almost tailor-made for this particular question that I'm going to be studying. Why? Because the sharp regression discontinuity design requires the assignment to the treatment condition to be exogenous and nonmanipulatable, right? I'm sure I'm butchering that word, so just bear with me.

So in terms of the used car market, think of what the treatment condition is. The treatment is whether the vehicle has expired manufacturer warranty or non-expired manufacturer warranty. And the decision -- your assignment rule to the treatment condition, which in this case is expired, is purely deterministic. So once you hit the age mark or you hit the mileage mark, you're in the treatment condition. Okay?

Second, it's completely exogenous. Why exogenous? Because it's predetermined -- even well

before it got out of the factories. And we're looking at used car markets, so we're looking at several years after these levels were set. However, going back to what I said on the previous slide, regression discontinuity design also affords us a nice way to control for these threats to identification, one of which is the role of unobservables.

And for some of us who are familiar with regression discontinuity design -- I see many in this room who have worked in this space -- by the choice of bandwidth, which is local region around the treatment condition, allows us to almost make the unobservables random -- as if it is random to the treatment assignment. Okay?

And I'm going to kind of try to provide some evidence and try to do as much convincing as possible with the data that I have that those conditions are being met.

AUDIENCE: Does the supply also vary around that cutoff, though?

DR. VENKATARAMAN: Supply of vehicles?

AUDIENCE: Supply of vehicles you have is --

DR. VENKATARAMAN: Yes, yes. And I'm going to find a way to convince you that that's not necessarily at work or that's not driving necessarily

337

the outcomes here. But it's a great point.

Okay. Validity tests. Remember, the threats to identification, I need to take into account the notion of sorting, manipulation, so all these things that are several tests that have been proposed in the literature to kind of allay some of these concerns, as I've been told and I've come to understand through the review process now that none of these tests are foolproof, which -- fine, right? However, if I can show a battery of tests, all of which reject these concerns, I'm hoping that this allays some of your concerns, right, otherwise, there's another journal, right?

So I have multiple editors here. I shouldn't be saying these things. Right?

But last but not the least, one of the limitations with this approach is, of course, external validity, right? Which is I can make a lot of fairly precise statements within the local region, and I'm going to refrain from making any statements outside this local region. Okay, so, if I say anything that's more preachy outside this region, call me out on that.

Okay, so the data set that I'm going to be using in this particular exercise, I got very lucky. I have 50 randomly chosen dealers across five states.

338

If you can -- if you have your geography right, these are states that border my state, North Carolina. And I have data spanning 2009 to 2014. For every vehicle that was sold at these dealerships, I have all the information.

So I know the purchase price, the list price, I know the profits that are realized from each sale, both from the sale of the vehicle as well as add-ons, as well as gains from extended warranties alone as well. But my NDA precludes me from exporting any of that information for the empirics or reporting any of those numbers. So my hands are somewhat tied.

For the purpose of the analysis, I'm going to limit all this analysis to B2C transactions. I'm going to be focusing on the top 15 make/model combinations which account for around 85 percent of all sales. Estimation sample boils down to 20K-odd observations. What I need is variation in the residual manufacturer warranties. So, as I said, there are those two types of manufacturer warranties, so here is a distribution of observations that I have across the entire mileage spectrum.

So I have 47 percent of the data resides in the region before the bumper-to-bumper warranty expires. I have 36 percent that resides between

339

bumper-to-bumper and a powertrain expiry. And then I have data past powertrain expiry. And that variation is what I'm going to be exploiting to the fullest.

So the first step to regression discontinuity design is going to be coming up with the local bandwidth. So we tried two approaches. Both of them seemed to be the de facto -- almost de facto standards in this particular area of research, one of which is the Imbens and Kalyanaraman paper and the Calonico Econometrica 2014 paper.

So I have multiple cutoffs. So I have the bumper-to-bumper; I have powertrain -- and powertrain could be either shorter powertrain or longer powertrain. So for each of those mileage markers I run -- I select my bandwidth, so I have a compact bandwidth around each of those markers. Once I had those markers, at this particular point, I'm going to be basically estimating -- running a regression. And this regression is going to be nothing but I'm trying to run logit -- not trying -- I'm running a logit transformation of the conditional choice probability. So it basically becomes a simple linear progression -- or nonlinear regression.

So what we have, the key parameters of interest for me, are going to be beta-one -- I should

340

walk you through the subscripts. So I is the consumer; J is the vehicle; D is the dealer; T is time. So D is an indicator variable that takes on value 1 if the vehicle is past manufacturer expiry for that particular mileage marker.

Notice that I have a slope in the region pre-manufacturer warranty expiry; and then I have a slope that is interacted with the indicator variable. So that allows me to pin down the variation in choice probabilities before and after. And the indicator variable is going to be allowing me to pin down the discontinuity at the point of expiry of the vehicle or manufacturer warranties. Okay.

How do I allay some of the threats to identification? So first of which is the strategic sorting between -- under composition of customers for one. Second, and this is a related point, and this is thanks to Matt for bringing this up, it's very possible one of the things that I don't observe in the data is the marketing effort by the dealerships to these individuals, right?

So suppose dealerships had perfect foresight and they had access to the analysis that we've undertaken. If I can show that the likelihood is higher pre-expiry and if the dealerships knew that,

341

perhaps they are more likely to market more aggressively to those people pre-expiry than otherwise.

So one way for that to manifest in the results or in the data should be if that were true we should see more bunching of observations in the pre-expiry than the post-expiry. So we should just -- we should see more observations before the expiry than after the expiry of the manufacturer warranty.

It so happens that the McCrary test -- not a foolproof test -- but it's one test that everyone's employed, that shows -- that allows you to test whether there is discontinuities in the density of their data. Okay, so that's what we employ.

Two, and it comes to endogenous selection of the marketing mix elements. So one thing that we do is we run regression discontinuity designs on all the continuous covariates that we have in our model. That allows me to assess whether there are departures to the left or to the right of the expiry of the manufacturer warranties.

And last but not the least, to ensure that I'm actually pinning down what I claim to be pinning down, I run a bunch of placebo tests. And what I mean by that is can I quantify or can I recover any

342

departures or discontinuities in my results in regions where I shouldn't be expecting any of these discontinuities. So we do it across several bins, and we are able to rule out the possibility of these departures or discontinuities happening anywhere else but where it's supposed to happen.

I'm hoping through -- and in the interest of time, I'm not going to walk you through the technical details of each of these things, but they're all in the paper. So the front end of the paper is actually much shorter than the back end of the paper because we have a lot more battery of tests than the actual formal main model itself.

Okay. So this is the regression discontinuity plot without any covariates. So the only indicator -- the only variable that we have on the right-hand side is whether it is pre or post-expiry of the vehicle, okay -- expiry of the manufacturer warranties. So as you can see, for the people in the back, you might not be able to see this clearly, but there is a very small discontinuity on the far left panel on the top, but there are a lot more noticeable departures on the other two panels. Of course, this is without covariates.

So clearly there could be other things that

343

could be driving choice probabilities. So once we include all these other covariates that I mentioned a few slides ago, instead of reporting the parameters, I'm going to simply highlight the -- graphically highlight the findings.

Okay, so this is what we find. So we see almost a linear increase in the likelihood of purchase of extended warranties leading up to the expiry of the manufacturer-backed -- manufacturer warranty, in this case the bumper-to-bumper warranties. And the point of departure or point of expiry of the manufacturer warranty -- in this case bumper-to-bumper -- we see a 3 percent drop in purchase rates. Then we see a constant attachment rate, a purchase rate, from that point onwards going forward. So if I were a dealership and I had access to this data, the first set of people that I'm going to be targeting are the folks who are between 35,200 and 36,000 [sic], because they had the highest likelihood of purchase.

The next possible candidates that I'm going to be targeting are any -- are all the individuals who are within my local bandwidth from the point of expiry of the manufacturer warranties all the way up to 48,000-odd-miles. The third best candidates are going to be folks south of 35,000. The further out you get,

344

the less attractive they become.

Tim?

AUDIENCE: What do the prices look like for these warranties over time? Is there a price discrimination feature associated with these?

DR. VENKATARAMAN: Yes. We have that. So we have prices in the model. So, most often, prices are increasing, as you get closer -- as we go closer to the extended warranty -- manufacturer warranty.

AUDIENCE: And what about after the --

DR. VENKATARAMAN: It's almost flat. You see step function. You see step functions, but the step -- yeah, you see step functions at 30,000 miles.

Okay. So, what might I be able to glean from this finding? Okay, so powertrain warranty we get just the opposite result. In the interest of time, I'm going to focus on what insight policy development as well as marketing relevant insight might I be able to glean from this result alone. So if those of us who are familiar with the warranty literature, the literature advances four mechanisms that drive warranty choice, as well as provisioning, first of which is insurance provisioning, right, insurance motive.

So think of this as an insurance. If I'm

345

risk-averse, I'm drawn to products that have more insurance. I'm more likely to purchase insurance products. So if the residual on my automobile is high, then I feel like that is less a risky product for me to commit to. Since I'm a risk-averse consumer, I'm going to be drawn to that particular product and but since I'm also risk-averse, I'm also more likely to purchase extended warranties.

So you should see more people committing to younger vehicles pre-expiry, and these very people are also more likely to purchase extended warranties. If it is signaling, think of signaling as the more insurance you have, it's almost equal to having higher quality product. If you have a better quality product, it reduces the need to purchase extended warranties. So the predictions from insurance motives and signaling are just the opposite.

Incentive motives, these have got nothing to do with consumers; it's got to do more with the firm side. So these have no bearing whatsoever on our results.

Sorting mechanism is the risk-averse consumers are going to be -- the more risk-averse you are, the younger the vehicle you're going to commit to; the less risk-averse you are, the more likely to

346

purchase older vehicles. Okay.

So what does this -- if these mechanisms were at work -- and I'm going to be done in a second -- what would they suggest? How might I be able to rationalize those two pictures? This would suggest at least as you're seeing an increase pre-expiry, that would suggest that insurance and sorting motives dominate in the region pre-expiry of the manufacturer-backed warranties when it comes to bumper-to-bumper warranty. However, when it comes to the region for the powertrain warranties, you find the opposite effect, in which case signaling motives are more at work.

So from a policy standpoint, this would suggest that in the region pre-expiry for the bumper-to-bumper warranties, the manufacturer-backed warranties and the extended warranties, at least in the minds of the consumers, are being treated -- almost traded off as complements, whereas in the region in the post-powertrain expiry, these two products seem to be treated more as substitutes.

So knowledge of these being as -- being either substitutes or complements is going to be directly pertinent to policymakers because based on that they can design interventions to either allay

347

concerns of -- concerns that need to be allayed, or kind of promote additional purchase of these particular products if it is actually economically prudent to do so.

Okay. That's pretty much all I have to say. So, happy to take any questions and refer to our discussant at this point. Thank you.

(Applause)

DR. JIN: The discussant is Matthew Jones from the Federal Trade Commission.

DR. JONES: Thanks. I have no slides. I just have a few brief comments, which is mostly because I think it's a very clean and straightforward application of an RD design. So not a whole lot to say, but I do have a few suggestions.

But, first, let me just review the punchline of the paper. The main question is, is there a systematic variation in the probability of purchasing an extended warranty around base warranty expiration. And the answer is yes. For the 36,000 mile bumper-to-bumper, the probability of purchase increases up to expiration, at which point there's a discontinuous drop. And then it's constant. And for the 60,000-mile powertrain, it's a constant probability, and then at expiration, there's a discontinuous jump, after

348

which it declines. That's the finding.

And the approach, I think, is a very nice application of RD design, given that there's no strategic variation in warranty expiration. So you might worry about -- well, you don't have to worry about manipulation of the mileage on the vehicle, right? It's illegal to tamper with an odometer, so that's not a concern.

If you're concerned about strategic offering for sale of vehicles, contingent on, you know, whether you're just before the expiration of the base warranty or just after that, there's a test for, you know, the density. And the finding is that there's no difference in density of offering -- or for up-sales on either side. So it seems to be a very clean implementation.

And, you know, the findings, I think there's an intuitive interpretation, which Sri just explained. So you have the sorting by risk aversion to explain the bumper-to-bumper, that more risk-averse consumers are more likely to buy a vehicle that still has a warranty, and also more likely to extend that compared to less risk-averse.

For the powertrain, the 60,000-mile powertrain, where reliability might be more of a concern, it's an older vehicle, there you might have a signaling thing. So the fact that the manufacturer still has a warranty on this vehicle effectively is a guarantee of quality and makes it less likely that this vehicle is going to break down. So I'm less concerned about buying an extended warranty. And I think that's an intuitive rationale for the opposite finding for the higher mileage warranty.

But just a couple of suggestions on the paper. So the statistical significance is brought out in these results. But I think it could be a little bit stronger in explaining the economic significance of the estimates. So if you look at the magnitudes for the bumper-to-bumper -- or, sorry -- yeah, bumper-to-bumper in particular, 36,000 miles, the 3 percent discontinuous drop is less than 1 percent in absolute terms. So less than 1 percent point change in the probability of purchase.

It's not obvious to me that that's economically significant in terms of motivating strategic targeting, if the effect is -- or if the curve is relatively flat. That's not to say that it isn't economically significant. But, I think, you know, it would be nice to know some more about in concrete terms about what does this mean for a manager.

These are tremendously profitable products, so there may, in fact, be evidence that it is economically significant, even if it's a small magnitude.

Also on the causality issue, I think there's one limitation, one thing. So if -- you know, an identifying assumption here is that all the covariates might otherwise explain a purchase are smooth around warranty expiration. There's one covariate that isn't measured.

And, you know, I think you've done everything you can within the limitations of your data to address these things. So that's one of the things about the paper, I mean, all the tests that you could do with the available data are done, and you get a result that confirms.

But there's one thing that isn't observed, and I think that's exactly how is the extended warranty presented to the consumer in the F&I office, right? So, you know, you go in, and if you think about how you would design an experiment, right? A consumer comes in to the F&I office, and the F&I manager presents a series of optional add-on products. And what you would want to have if it was a controlled experiment is you'd want to have those products presented in the same way on either side of base

warranty expiration.

So one way this could go is the F&I manager says, you know, here is one out of 30 pages you have to sign. This happens to be the extended warranty; it costs $2,000; do you want to buy it or not. And the consumer just responds, right? And it's presented the same way whether or not there is a base warranty.

Another way that it could be presented, it could introduce a framing effect where, you know, they say either your vehicle still has a warranty on it but it almost -- it's almost expired, you might want to extend it, it costs $2,000. But if you're on the other side of base warranty expiration, your vehicle does not have a warranty; would you like to purchase one? And I think something like that, while you can't observe it, could, you know, produce an effect such as a discontinuous jump at expiration. So that's just one possible limitation.

But overall I think it's a very nicely executed and interesting piece, and it's encouraging to see evidence that consumers are responding to real economic incentives and information in this decision -- in this purchase decision, rather than just sales pressure, which I think is something that is adequately tested for. It's just the framing of

the sales pitch may differ in a meaningful way.

And that's all I have.

(Applause)

DR. JIN: Thank you, Matthew. We can take a few questions.

Sri, you want to come up?

DR. VENKATARAMAN: Sure.

DR. MISRA: Just maybe a thought about this. So at 36,000 miles limit, right, so the drop, now -- so there is a sorting argument based on risk aversion, but -- and maybe this does happen, that firms obviously sometimes might have incentives to pre-announce certain kinds of incentive schedules such that they actually preselect everybody before the threshold -- the ones that have to buy, and for the ones that are left behind are the ones who have been kind of endogenously selected for -- so this could be another kind of sorting which probably might be optimal for firms.

DR. VENKATARAMAN: Beautiful, Kanish, great point. So the only observations that we are tackling in this particular analysis are the sales that are consummated at the dealership, right? So correct me if I'm wrong, what I could do -- in your setting, you're possibly talking about sales or offerings that are presented to consumers post-purchase of the vehicle. So you're sitting at home and you receive these mailers. So perhaps some people were strategically chosen to receive it and possibly even the framing of the message was slightly different.

DR. MISRA: Who would receive actually mails from your dealer post --

DR. VENKATARAMAN: Yeah, and I can tell you having spoken to F&I people and underwriters, they do blanket mailing. Everyone receives it, right? Some markets, what they do is they receive multiple messages from the same individual, and the only thing that I've been told that they change is they make the reminder note, sometimes you have seen in the picture as well, this is the last reminder.

Apparently, there are some people -- some of our peers who seem to view that as, you know, with a greater sense of urgency when someone says this offer is going to end tomorrow, and they feel like they're going to lose out on something big, and they commit to these products. But great suggestion.

DR. JOHNSON: One question in addition as well. Like you were mentioning about sorting. And I was thinking, like, one way to probably tease out a little bit more of the effect might be the gradiation and

behaviors across different types of models which are different in terms of their reliability, right? So did you explore that at all?

DR. VENKATARAMAN: Yeah, so I tried to condition it by having the model fixed effects, right? So one of the things I could do --

DR. JOHNSON: But then you're just absorbing everything.

DR. VENKATARAMAN: Yeah. So one thing I could do is get that out and kind of recover the treatment effects across all make/models. We tried to do that. There are a whole bunch of really unreliable vehicles for whom we don't have much data. So much of the action is actually coming from parsing of the variation and the slightly more reliable data.

So remember, the more reliable vehicles are also the more expensive vehicles. The more expensive the vehicles, the more expensive it is to repair anything if something were to happen. So we're actually -- much of our identification is actually coming from the higher end set of vehicles. But great observation, great intuition.

AUDIENCE: So I might have missed this. I was just wondering, you know, Garrett's question a little about supply. I missed what kind of --

DR. VENKATARAMAN: Yeah, so in one of the -- one of the things that we do is we try to assess if the supply of similar vehicles in the local market around this particular dealership has any bearing on the likelihood of purchase of extended warranties. And the thinking is as follows. If you buy really old vehicles, the supply of parts for these vehicles is a stock of existing vehicles in that local market that are going to be, you know, turned in as salvage vehicles.

So in order to proxy for that effect, we kind of include the stock of variables of similar type, similar age, similar vintage in that local market. We do have many of that particular vehicle, type of vehicle in that market. We just don't have it in the back lot of this particular dealership. And we kind of test whether that has any way to explain some of this variation, that variable to pick up.

DR. JIN: I wonder what role price plays in this whole thing. For example, would the dealer lower the price of the car in order to persuade the buyer to buy extended warranty?

DR. VENKATARAMAN: Right. So, we -- yeah. So we actually explore that to the fullest. So what we do is we actually run a regression discontinuity on

the transacted value of the vehicle and try to see whether the prices, all else being equal, are systematically lower for that kind of vehicle, pre versus post, and we don't see any difference. We rule that out.

DR. JIN: Which is surprising.

DR. VENKATARAMAN: Apparently, the outcome, at least what was told to me, is the F&I people are compensated for the extended warranties; the sales guy is compensated for consummating the business deal, so -- which might explain why those two are not necessarily going hand in hand.

AUDIENCE: What about certified pre-owned? Sort of a combination between --

DR. VENKATARAMAN: Great.

AUDIENCE: Do you exploit that at all, or --

DR. VENKATARAMAN: We don't.

AUDIENCE: -- see those --

DR. VENKATARAMAN: We don't. The reason we don't is because most certified vehicles, what it does is the dealership or whoever is certifying the vehicle, you're basically extending the manufacturer warranty. So in the data, all we know is this particular vehicle has been certified. I just don't know how long the warranty has been -- has been

357

extended. So in order to kind of mitigate any issues that might arise as a result of those observations, we kind of keep them away from the analysis.

Any other questions?

AUDIENCE: Yeah, I was just wondering if your study looked at whether there were any massive recalls during the time that you looked at and that were, you know, widely publicized and whether that had any impact on a consumer's decision to purchase an extended warranty.

DR. VENKATARAMAN: During the period of our data, we had four major recalls, and we have -- in one of the specifications that we tried, we actually had recall indicator variables for those make/models. I don't remember off the top of my head what we found, but all I know is we decided not to put that in, largely because for the most part it wasn't explaining any of the choice proclivity.

Remember, this is pre the big Toyota recall. But thanks for that observation. Thank you.

(Applause)

358

WHAT DETERMINES CONSUMER COMPLAINING BEHAVIOR

DR. JIN: Thank you. We'll continue with the last paper about consumer complaining behavior represented by Devesh Raval from Federal Trade Commission.

DR. RAVAL: Thanks.

So thank you all for staying until the last paper. I know that I'm -- you guys have better things to do, but I'm glad that you're here. I want to thank the organizers for inviting me and also the people that have been laboring behind the scenes. And I'd also like to thank Anne Miles and Patti Poss, both of whom are by the windows for helping provide the data and also asking me lots of pesky questions that helped develop the paper.

So let's start. First is the obligatory disclaimer. What we're interested in in this paper is about consumer voice. We have known at least since the work of Al Hirschman that consumer voice really matters for markets, but the amount that it matters has increased a lot since the advent of the internet. So essentially through things like user reviews, the internet has allowed the effect of consumer voice to be magnified.

So there's been a lot of work, both in

359

economics and marketing, showing that user reviews affect demand and, in fact, there's a recent paper by Greg Lewis showing that the amount that user reviews affect demand has increased a lot over time, which is consistent with the idea that, as consumers get more comfortable with the internet, they also use the internet more as a source of information.

Now, consumer voice can affect a lot of different decisions. So for a consumer, this might be just a measure of product quality. So I'm thinking about going to a restaurant; I look at Yelp to see how good it is. For a platform or a retailer, it may tell you which products to display or stock. So Yelp is going to look at the reviews and then pick things that have high star ratings to show you. And then finally for the manufacturer, it's going to show you how to improve products.

So imagine I'm a restaurant and, you know, there's a principal agent problem, I don't know what my staff is doing, I can look at Yelp, the people who are complaining about, say, the waiters or they're complaining about the food, then I know what I need to do to improve. So really this kind of consumer voice can help to improve a lot of different dimensions of

360

the market.

But there's still a question of whose voice do we here. And, so, in general, we know very little about the characteristics of reviewers, and it's likely there's a lot of self-selection. So, you know, for example, I've been using Yelp and Amazon for a long time, but I've never written a review. And I imagine only a certain fraction of you do write such reviews.

Self-selection could affect a bunch of different parts of this. It could affect which products are reviewed, as well as how quality is assessed. So I think it's easiest to understand this through a set of examples. I have a couple of examples here.

So the first one is to think about franchise hotels. So I am a big hotel chain. I want to know which of my franchisees are, say, performing service adequate to my brand image and which are not. Now, if consumers vary in their complaint propensity, then just looking at the sort of complaints or reviews I receive might be very misleading. So for example, one of the main messages of this paper is that consumers in heavily minority areas complain less.

You know, so if I have some hotels that are

361

being served -- that are serving mostly white customers, others are serving mostly minority customers, it could be that the ones serving mostly minority customers look better than they actually are. And that's because those consumers are not willing to review.

And, so, you know, from the franchise hotel's perspective, it's hard for them to know who are the good managers, who are the bad managers, which is the good franchise, which is the bad franchises. From a customer's point of view, the sort of reviews they see online may not provide a good estimate of quality.

Now, the second one I've given here is the Consumer Review Fairness Act. So this is actually recently passed by The House, I think in the past week or two, and what they're trying to do is prevent firms from penalizing people from making complaints online.

So if you think about it, if firms are allowed to penalize people that make complaints by threatening to -- by threatening to fine them or something like that, then you might have a lot of selection where the left tail in the distribution is not being voiced because people are afraid that, you know, there will be retaliation if they say something.

362

So this might be things like, literally, you know, you get sued or you get fined, but I talked to Steve Tadelis when he was chief economist of eBay, and one of the things he was saying is that in general the review stars of buyers and sellers were very uninformative. And the reason is that there's retaliation if I give a seller a one-star rating, then they're going to give me a one-star rating. And so the better sort of informative signal is whether -- you know, what the fraction was of reviews you get rather than the actual star rating.

Now, in general, there's kind of a fundamental identification problem here if you don't have consumer experience data. And that's because if you see higher rates of consumer complaints, that could be because those consumers have a higher propensity to complain, or it could be because they have a worse consumer experience. And, in general, it's going to be very difficult to disentangle these two stories because in general we don't have this kind of consumer experience data unless maybe you're a big internet firm that knows who purchased all the products.

So I'm able to separate the two stories using a set of legal cases. And Patti again helped

363

provide some of these cases, or at least told me which I could use. But right now we have a bunch of victim data sets that are matched to complaints from Consumer Sentinel Network.

So I don't know if you guys know what the Consumer Sentinel Network is, but it's an organization that's getting complaints both from a lot of government agencies like the Federal Trade Commission and the Consumer Financial Protection Bureau, but also private actors like the Better Business Bureau, which receives millions of complaints a year.

So in this, we have a bunch of cases where we have a data set, sort of the customer data set of a company, that's all the victims of a particular scam, and then we have matched all the complaints about that company that we were able to get from Consumer Sentinel Network.

Now, what's crucial here is that we have addresses, in general both for the victim data sets and the complaints, and that means we can link these to demographics at the zip code level. And there's actually a very important policy question here, which is, you know, one of the things -- one of the ways we use the Consumer Sentinel Network, it's called a sentinel, and the reason it's called a sentinel is

364

we're trying to look forward to try to, you know, identify emerging problems in the marketplace and try to solve them before too many people get victimized.

And, so, we want to make sure that we learn about problems affecting all communities. And if certain communities are a lot more likely to complain than others, then we might just be responding to the problems of one group of society and not other groups. So this is a big policy question for us at the FTC, as well as places like the CFPB.

So let me go over the main takeaways quickly of the paper. So what I find is that there's substantial selection in complaints. So areas with more minorities, the areas with more blacks and Hispanics complain at lower rates, whereas areas of more college graduates complain more.

And, crucially, it's really important to control for consumer experience. So if you just compared complaint rates for population, what you're going to find is that heavily black areas complain about the same or maybe more. So you're going to get a very misleading picture of what's going on in the marketplace. And that's because some of those heavily black areas are going to be more -- victimized at higher rates, and so they're complaining more, even if their underlying propensity to complain is lower.

So let me talk a little bit about the related literature of this paper. So this is sort of in between the literature on customer reviews and the literature on customer satisfaction. So as I said, there's been a lot of work showing that customer reviews affect demand. I've highlighted three other papers. So the first paper Dina was talking about at the panel, which is that there's strategic behavior going on, so, you know, somebody may write false negative reviews of their competitors; they might write false positive reviews of themself.

Second, Ginger has a paper on how to optimally rank given reviews. So if reviewers vary in their mean in variance and other characteristics, you can use that to provide a better ranking than just the average star ranking.

And, finally, there's been a little bit of work on how reviewers or reviewer characteristics demand. Now, second, there's been a large literature on customer satisfaction, and there's even a journal dedicated to customer satisfaction, as Jan has pointed out to me multiple times. But the foundation of this literature is from the book Exit Voice and Loyalty by Hirschman. So he sort of started out the theory of

this. And then there's been a large empirical literature.

But, in general, this empirical literature hasn't really been satisfied -- satisfying. There's not really been a consensus about how different demographic groups vary in their complaint rates. And there's two reasons for this. First of all, in general, the samples have been small.

So the sample might be, you know, a thousand-person survey, or it might be data from one local Better Business Bureau. But I think more importantly there's been no control for consumer experience. So you don't really know what's being identified when you just run these regressions, which is part of what I'm trying to do in this paper.

So let me go over the demographics first I look at. So I'm matching consumer zip code to ACS 2008-2012 demographics. So I look at a bunch of different demographics, and I'm going to be controlling for all of these, but in the paper -- in the presentation, I'm going to focus on percent black, percent Hispanic, and percent college graduates. I also look at urbanization, median household income, unemployment rate, median age, household size.

So this is getting, you know, a bunch of

different attributes of different groups. And, again, this is all at the zip code level. And I also try to discretize all the demographics in order to allow nonlinear effects of demographics. So, for example, really high-income areas might be not any different than low-income areas.

And as I said, I have four cases where I can compare victims to complaints. So it's kind of nice in a paper to have a little bit of mystery. So here the mystery is that the first case I can't tell you anything about, so I've called it Case B, and it's nice because it has over 12 million victims and over 4,000 complaints, and it's by far the biggest data set. The only thing I can tell you about it is that it's been successfully sued in court by a state or federal agency. So that's -- you know.

Now, the second case here is Ideal Financial, so this was an FTC case, and what the company did is bought payday loan applications and then withdrew money from the bank accounts of people. So if you do a payday loan application, you have to give data about yourself, and then they just took money from those people. So here we have 2 million victims and about 1,500 complaints.

The third case is Platinum Trust. This is

also a payday loan-related case. So here they took payday loan applicants and they tried to sell them deceptive credit cards, so credit cards that weren't real credit cards but maybe they claimed they were. So this is the smallest data set. We've got about 70,000 victims and 500 complaints.

And then the last case, WinFixer, is a spyware case. So this company falsely claimed that you had spyware, and then it was going to sell you software to then, you know, remove the spyware. So we've got 300,000 victims and 1,000 complaints.

But in general, these cases are all somewhat different from each other. But I think the key thing to note is that the number of complaints is orders of magnitude smaller than the number of victims. So this is telling you that there's a lot of selection in who decides to complain. The average person is not complaining.

So the first thing I want to show you is that complaint rates are correlated with victim rates. So here I've just constructed the complaint rate per capita and the victim rate per capita and the zip code level. So as I showed you here, the number of complaints is pretty small, so they're about 25-, 30,000 zip codes in the U.S. You know, even the biggest case is about 4,000 complaints. So the complaint rate is usually going to be either zero or, you know, there's going to be one complaint or zero complaints divided by the population. The victim rates are going to be much bigger because we've got millions of victims for a lot of the cases.

What I do is I try to provide a standardized estimate of, you know, if you increase the victim rate by a standard deviation, what happens to the complaint rate. And what I find is that across these cases we find pretty significant effects. So areas with higher rates of victims also have higher rates of complaints.

And the magnitudes are about the same across cases. So it's about -- if you increase the victim rate by one standard deviation, the complaint rate is rising by 12 to 17 percent. So this is actually pretty reassuring. It says that the data isn't crazy. Areas that have more victims in general are going to have more complaining consumers.

Now I want to look at demographics. So the first thing you see is that the victim rates across demographics vary widely by case. So the X axis here is the population share that's percent black across zip codes; the Y axis is the victims per thousand population. It's just normalized by the mean, so

everything fits on the same graph.

So the blue and the green line are both the cases that involved payday loan victims. And what you see is that when you go to areas that are 100 percent black, they have complaint -- they have victim rates of about 300 to 500 percent greater than areas that are 0 percent black. And my guess is that that has to do with the kinds of consumers that buy payday loans.

The Case B that I can't tell you about has higher rates of victims in heavily minority areas or heavily black areas, but it's only 80 to 100 percent larger. So it's not quite as much as those two cases. And then the WinFixer, the spyware case, is about flat. So areas that have very high percent share blacks have about the same victim rate as low percentage share blacks.

If I look at Hispanics, things look more similar. You see sort of an inverse U shape, so it seems like the highest victim rates are in sort of moderate, 25 to 50 percent, Hispanic areas. We find lower victim rates in really high Hispanic areas. And -- but in general the variation here is not quite as large as it was for percent black.

So here I'm showing you that if you look at here, the huge differences across cases and victim

rates, but if you look at the complaint rate -- so this is the number of complaints per thousand victims, again normalized the same way. What you find is across all four cases, you see a decline. So areas with high population of blacks have -- have less complaints relative to the number of victims than areas with low percentage share blacks. So this is about a 40 to 80 percent decline.

So this is just the raw data. So I want to copy out this a little bit in that, you know, again, most of the zip codes have zero complaints because we don't have many complaints. And, so, this is nonparametric, but I think if you did a statistical test, it's going to be hard to get non -- strong evidence with this kind of data, but I think this shows you that if you just look at the raw data you do find that heavily minority areas complain less. If you do that with Hispanics, you see the same sort of pattern. And, again, heavily Hispanic areas that are close to 100 percent Hispanic have about 50 to 90 percent lower complaint rates than areas that are 0 percent Hispanic.

So I'm going to then examine this more formally by using an order logit on the individual data. So here the Y variable is a latent variable for

the demographic category. So what I've done is I've discretized the categories. So the five categories for black and Hispanic are sort of 5 percent black, up to 25 percent, 25 to 50, 50 to 75, 75 to 100. Now, the main coefficient of interest is alpha, which is an indicator of whether it's in the complaint data set versus the victim data set. So what I've done here is I've just stacked the different data sets together, and this is trying to compare how do the demographics vary between the complaint data set and the victim data set for a particular case.

And then I've put in controls for all the other demographics, as well as the log population. So what this is trying to say is if you go between the victim data set and the complaint data set, how does one particular demographic category vary, even after you control for all the others.

So I'm going to point out that because I'm using the ordered logit, that's putting a lot of structure, so that's going to say that there's going to be -- if you go across demographic categories, that's going to be a particular downward or upward pattern. And you need this kind of structure because at the end of the day we don't have that many complaints, which is part of the self-selection

373

problem to begin with.

So here I've graphed the confidence intervals for the percent change across categories when you go to the complaint data relative to the victim data. So it's easier just to look at the bottom right corner. So the bottom right corner is areas that are 75 to 100 percent black. What you find is that you find significant negative percent changes in the complaint data across all four cases, and this is about a 25 to 80 percent decrease in complaints relative to victims depending on the case. So this is blacks.

If you look at Hispanics, you see a similar picture in the sense that for all four cases you find declines, significant in three of the cases. And these are actually bunched pretty close together at about a 30 to 40 percent fall in the complaint rate for very heavily Hispanic areas.

Now if you look at college-educated areas, you find higher complaint rates. So this is about -- if you look at the areas with greater than 60 percent college graduates, you have about 25 to 50 percent higher complaints relative to victims. So the paper has all the other demographic categories, but I didn't want to bore you too much. So I've talked about the

374

ones that I think are the most interesting.

So what this says is that there's a lot of self-selection. Heavily black areas and heavily Hispanic areas complain a lot less. Heavily college-educated areas complain a lot more. Now, you get very different patterns if you don't control for customer experience. So this is what the literature has done in the past, and it's sort of a more naive thing where you're going to look at per-capita complaint rates and see how they vary with demographics.

So here I take data from the Consumer Sentinel Network from 2012 to 2015. I exclude identity theft data. And the specification here is I'm going to look at the log of the expectation of the complaint rate as a function of demographics, population, and time and state trends. So first let me just show you the nonparametric regression. So, again, here, the X axis is the population share of percent black or percent Hispanic. The Y axis is the number of complainants per thousand people.

So for percent black, you see it's pretty flat. So really low percentage black areas complain about the same rate as really high percentage black areas. This is very different than what we saw when we looked at the cases where we could match complaints

375

to victims. So what this says is that, you know, if you don't have that kind of victimization data, you might have a very misleading picture of what's going on. The percentage Hispanic does decline just as we saw in the victim data.

And then you can do this formally and econometrically with that specification. What you find is that the red here is confidence intervals for the entire data set; blue is for the FTC; and green is for the CFPB. So the Y axis is the percent change in the complaint rate. And then we've got four groups, so everything here is relative to 0 to 5 percent black areas. And what you see is that for the entire data and for the FTC, heavily black areas complained a little bit more once you control for all the demographics.

So if you don't control for demographics as I showed you in the nonparametric regression, it's about flat. When you control for all these demographics, you find somewhat higher complaint rates for heavily percentage black areas.

And the green CFPB, you'll see huge increases. So, you know, for -- in the CFPB data, heavily black areas are complaining about 100 percent

376

or more. And I think some of that is that they're complaining about -- heavily black areas are complaining about different things. So let me just show you that quickly. So it's hard to see all these different colors, so I'll just try to summarize what's going on.

And we're looking here at the percent change in the share of complaints where I've divided complaints into different categories like auto-related complaints, imposter complaints, debt collection, et cetera. And what you find is that in heavily percentage black areas, you get a lot more complaints on things like banks, debt collection, and auto-related complaints.

And I think there's a common theme across all of these, which is finance because I suspect a lot of the auto-related complaints may be related to auto finance. So what this is saying is that heavily black areas are complaining about different issues, and likely that's due to different rates of victimization or things like that.

So I guess I have a couple minutes left, and so I'm going to talk about, you know, what can you do with all of this. So how should we account for selection. So I think there are two potential

377

answers, and I think there's some complementarity between those answers. So first of all, there's a policy answer, which is sort of outreach. So here we contact groups that typically complain less.

So for the FTC, this is things like outreach events, which we do periodically. So we might go to Atlanta or LA and try to hold an event where we talk to local community groups. We might want to talk to, say, non-English-speaking media, and try to get -- you know, first of all, tell them about the FTC, what we do, how they can complain, but also learn from them what their problems are.

Now if you're a marketer, this might be something like surveys or incentives. So you could think about running a survey of everyone that's bought your product and, you know, offer them a $50 gift card and then try to see what their -- see what their comments are. And that might give you a very different picture than just looking at the people that decided to review or decided to complain on the website.

And, again, incentives might be some way to get people to complain. So for example, you offer them a raffle ticket, essentially, to complain. And there's also a statistical answer, which is weighting,

378

so you could think about overweighting complaints from groups that complain less, but the problem with that is you need data on consumer experience to construct the weights to begin with.

So that's something, you know, I could do with this type of data because I have that data, but if you're a marketer, you might need to do sort of a survey or do something like that in order to do the weighting in the first place. But I think -- you know, I've not seen anyone do this in practice, but these are the sorts of things you would need to do in order to deal with self-selection.

So that's it.

(Applause)

DR. JIN: Thank you. The discussant is Anne Coughlan from Northwestern.

DR. COUGHLAN: Well, thanks very much. I feel I have great power, and yet you have great coercive power against me if I go long, so I'm going to go short.

But thank you for all -- yeah, I like to walk around. Better if I can walk around.

So thanks again for this great conference today. I think we've all had a wonderful array of papers and presenters and great discussions, so thanks

379

to the organizers and thanks for hosting to the FTC. And it's a pleasure to be here.

So I have some comments on Devesh's paper here, which I found very interesting. And, in fact, his presentation helped me a lot. So one of the first comments for Devesh is going to be a little more clarity in writing, please. So you'll see that one of the things that I have to say I think really was due to the fact that I was a little confused about what was being intermixed where in the paper, but we will get to that.

I want to emphasize to those of you who are not at the FTC what I found immensely interesting and novel here. There is a novel data set that is available only in the law enforcement community, which is the Consumer Sentinel Network. Now, for those of us not in the FTC, this reminds me greatly of the explosion of papers that I see of academic colleagues with people who work at Google and Microsoft, right, where you really cannot get the quality of data, right, sitting outside of the community.

So this is terrific because it compiles lots and lots and lots of different kinds of consumer complaints and very wonderfully it frequently, apparently, includes the complainer's address, thus

380

facilitating the entire analysis that's done here that lets you put together zip code data with -- on complainers with the actual nature of the complaints.

And, then, there's analysis that combines the demographics of the victims of fraud from this interesting sample of four law enforcement cases and asks whether the propensity to complain correlates with a number and type of victim.

So what I saw as the goal of this paper is a much deeper descriptive dive than I've seen before into the nature of who complains, which is super important for us to understand and also show how the demographics of complainers compares to the demographics of victims, which is very important potentially for consumers protection issues that we're concerned with here today.

I'm not going to go over again the interesting particulars of the data. You can take a look at this yourself. The interesting thing that I found here is I thought about this -- I thought about the question of why. There's a lot of interesting information here about what, and I was trying to figure out about why, okay? And, in fact, you see some of this in the paper.

One of the questions is is this due to

381

differences in people's cost of time? Is it due to differences in people's access to the ability to complain or the knowledge about how to complain? Because the answers to those questions are crucial for helping to get voice out there properly. So, for example, one of the things that wasn't emphasized in presentation but which I found very intuitive is that complaint rates are lower for areas that have higher household size.

Well, a few of us were talking about being parents of kids, and you know what that one's about, who has time to follow up on complaints when you hardly have time for four cycles of REM sleep per night, right? So that one was very intuitive to me. And the one that I found kind of intriguingly different from what I thought would happen is that complaint rates are higher in areas with a high percentage of college grads.

Again, if what you believed was the cost-of-time hypothesis, you'd guess this isn't happening, right? So I found some of these actually just very interesting on a univariate analysis, right? There's some very intriguing descriptives here.

Now, going on to the law enforcement actions, there are four different law enforcement

382

actions here, and as I saw it, and I believe that's the way it was presented here, too, relative to the level of victimization, if you're in highly black or highly Hispanic areas, you see fewer complaints. But with higher college education, you see more complaints.

And, so, the sort of inference that I wanted to draw from this is this, that perhaps the types of complaints in these sub-populations don't reflect the types of victimization of concern in the cases that were presented here. And I think there was a sense of that, that there's lots that people complain about. So is it really just that there's a larger array of issues to complain about in some areas, right? And some of this doesn't really concern victimization due to fraud.

So I believe some of this was cleared up, but it wasn't clear when I first read the paper whether the complaints database was restricted to fraud complaints. Now I believe from having heard the presentation that that's carefully culled down, but if not, you want to separate out and make sure you have fraud/fraud in the four cases looked at.

Now, one of the things I thought about was the classic alternative explanations, you know, string

383

of contemplation. So one of the thoughts that I had is this: What is, so to speak, an equilibrium complaining process? When is it that you would decide, so to speak, that on the margin it just isn't worth it to you to complain about whatever it is that is happening? And in particular, so many people do not complain, and some of these things -- payday loan frauds and so on -- presumably would be notable enough you would expect a lot of people to complain, and yet they don't all. Right?

So why do people not complain? That was kind of interesting to me. And perhaps on the margin what we want to think about is a sort of an economic model where on the margin the necessary number of complainers to sort of induce action, right, is really of interest here. It doesn't take, you know, however many hundreds of thousands were harmed for action to occur.

And, so, in some sense, I was thinking, well, perhaps this is really an okay number of complaints. I mean, we don't actually know what the right number of complaints is, do we? Right? The right number of complaints is the number that induces action to occur, and then that brings in my mind another thought, which is -- you may be familiar with

384

this -- the commentary -- I'm not sure how much research was literally done on this -- about what happens when violent crime occurs on the city streets. And when there are only one or two people who observe it, they tend to run and help. And you remember the famous case -- I can no longer remember the name of it -- I think it's -- Kitty Genovese, exactly. And how many people, 30, 50 people later on reported that they had heard about her being attacked, and she was killed, correct?

AUDIENCE: I believe that story is no longer true.

DR. COUGHLAN: It is not?

AUDIENCE: No.

AUDIENCE: Apparently it was made up by The New York Times reporter.

AUDIENCE: Yeah.

AUDIENCE: But it's a good story anyway.

DR. COUGHLAN: Well, okay. Well, let me state this as a hypothesis, then.

(Audience comments off microphone.)

DR. COUGHLAN: Thank you. Thank you.

So the hypothesis, though, is that if you know that there are many others, your individual impetus to complain goes down, right? And, so, it may not really be necessary for us to be seeking more and more complaints, and it could be helpful to understand to do some investigation into, well, what is the necessary number of complaints. Okay?

Now, the other thing that I thought I'd say a couple of words about, and then I'll close off, is some more ideas for continued research. We have an intermixture of four different cases here. Two of them are payday loan; one is about spyware; and the other one is -- I don't know what. But the two payday loan ones are similar, and the other two -- well, one of the other, the spyware one, is obviously very different; and the fourth looks different as well.

So one of the things I am thinking, and I know how burdensome it must have been to create the information per case, so I'm in dreamland. I'm not worried about the cost of data. But it would be interesting and probably valuable to cluster together like types of cases because then you could pool data and think about the common issues here, but it probably isn't appropriate to pool across these four because those are very different drivers for those cases.

And then there are lots of different types of complaints. And you saw some of that in the, you

know, auto and bank -- the categories of products, but there are also different types of complaints. There were customer service complaints; there were "I was overpriced" complaints; there were "I couldn't return my product" complaints; all kinds of complaints out there.

And, so, I'm thinking that conceivably this complaining kind of research, whether you do or don't want to go beyond consumer fraud, per se, could produce a whole wealth of interesting research projects, okay?

Finally, I've got this just as one small comment, but I don't know how possible it is, but it would be very useful to try to figure out metrics for filtering out spurious complaints because there are complaints that are not real, and that's an important thing to do, too.

So in sum, what I found so interesting and quite different from data available that I've seen in other projects is that here there's a degree of detail you just can't find anywhere else. So I would urge keep on going. There's a ton of interesting stuff here, with the possibility of getting a much better judgment on when and where you want to take action on consumer fraud.

So I'll stop with that. Thank you.

(Applause)

DR. JIN: Thank you, Anne. Any question?

DR. RAVAL: Can I give a response to one thing?

So I just want to give a quick response to the comments on this slide, actually, which is, you know, to try to think about more detail in what people are complaining about and not just aggregating complaints. So, in general, this is, I think, a machine-learning or text-finding challenge.

So we have the complaints; we have a categorization that it's about autos or it's about debt collection or something like that. But to go deeper, you really have to look at the text of the complaints, and I've done some work on that internally, but, you know, we have the free-form text of what people say, and there are some potentially crazy complaints. There are going to be people that complain about each one of the issues you talked about, but the question is, I think, how can you use something like topic modeling and machine-learning to try to do that.

DR. COUGHLAN: Exactly. I would totally agree. Yeah, it's not easy to do.

AUDIENCE: So about the machine-learning part on complaining, I think there's some studies, so using Facebook data they can pretty much see how many complaints, and Facebook has a natural policy. They changed the conversation of how a customer can complain, airlines, hotels, so that before the changes the two are not classed together.

So after the policy change, all the complaints that you complain about the surveys was bad, and then they cluster together. So there's a very good -- and that your policy can be wrong. You're really tied with your complaint behavior here with social media and machine-learning techniques. So we can talk more offline.

DR. RAVAL: No, that sounds like something we should be doing internally.

(Applause)

DR. JIN: So thank you all. Thank you to all the presenters and discussants and actually the active participation at the whole conference. It's really sweet.

So before Sudhir delivers the closing remarks, I just want to mention a few logistics things. One is that transcript of this whole conference will be available soon on our website. So

389

if you missed part of that, you will be able to get back to it in the transcript. We're also planning to post the slides on our website, and before doing that, we're going to email the presenters and discussants and make sure that you -- if you want to put some updates into the slides and you will be able to do so.

If you have any comments or suggestions about this conference or future activities, then we can organize with your marketing community or even now, other communities; you're welcome to send us an email at the marketingconf@ftc.gov, which is the same email website that you will see in the registration website. That's where we'll welcome your comments.

And, finally, I want to thank Laura Kmitch and Constance Herasingh for really running the whole show for the whole day. They not only made sure the computer worked, made sure the lunch worked, made sure the time worked, and the microphone worked, they actually have been helping me from day one, from planning to all of the probably followup work after today. So let's give a round of applause to both of them.

(Applause)

390

CONCLUSION/CLOSING REMARKS

DR. JIN: So with that, I'm going to turn over the podium to Sudhir.

DR. SUDHIR: So, first, let me start by thanking Ginger. As she mentioned, you know, I think just -- I think, if I recall, it was November of 2015 I was just taking over as Marketing Science editor and Ginger was -- I saw on LinkedIn that she was taking over the Director of the Bureau of Economic Analysis. And I sent her a note on LinkedIn saying, congratulations; by the way, we should do something with marketing.

And I really didn't have any clear idea what I was thinking. And a couple of months later, I get this very detailed proposal to me and Avi saying, hey, you know, we should put together a special issue. And I was just -- I mean, to me, I was thinking about really a special issue and, like, you know, this just seemed to be the ideal special issue, I think, to do. Because partly I think there is a fair amount of latent interest, as was evident from over the 100 people who registered and came to this conference.

And, so, my sense of it was that there's always this interest in the ability to do work related to consumer protection in the marketing-economic

391

community, but nobody had role models of papers that were actually targeted towards that -- those set of issues. And, so, we would always have some throwaway lines at the end of a conclusion saying this work would be relevant to policy regulators but with really no specific, you know, particular analysis that was done, a counterfactual that was particularly run or even some dicing and slicing of the data in ways that would be particularly relevant.

And I was reminded of that partly when Catherine Tucker was telling Hema, you should actually slice the data and take a very limited slice of data and see whether that would already do things other than privacy -- it would be nice from privacy perspective, you don't have to have a long data set. And as she said that, I was reminded of a paper that I was writing around a very similar issue that Hema was discussing, but it's sort of ad-targeting on price targeting. And one of the counterfactuals that we were doing, what would you do with last visit, last purchase; then Catalina, the company that did it, would keep only 64 weeks of data, and we had 100-plus weeks of data.

And, so, we wrote something with 64 weeks of data, but our motivation, sadly enough, now that I

392

think about it, was, you know, data storage is very expensive. Companies don't want to store data. Therefore, you know, let's try what we can do with 64 weeks of data. And we -- you know, and we then said, you know, storage is getting cheap; why the hell would you care about this, remove all the stuff from the paper.

So I was looking back at the paper today, and as you commented, and I found that we did not have the 64-week description, but if I had written, hey, given privacy concerns, if we had gotten 64 weeks, wouldn't it be wonderful, and everybody would have said we were so far ahead in terms of thinking about this issue. But it was a counterfactual that we had run, but we motivated based on storage cost, which made no intuitive sense to anybody.

But my point is that I think there is lots and lots of opportunity if you start doing the data with exactly the same kind of things that we would, but we would be informing people.

So, in fact, when Ginger sent us this thing for the special issue, Avi and I talked about it, and one of the things that we said was we should have a conference before we run the special issue because we wanted people to have a melding of the minds, so to

393

speak, talking to the folks and they're busy understanding what it is that they do, and also by listening to everybody's talks, think about what kind of questions would be interesting and inspiring because I really think a lot of us are already doing a lot of things, but we've just not been slicing the data as I just told you with my silly example, right?

So -- and that's why we have the deadline for the special issue as July 31st, another nine months, which means that any of the work that you're doing, if you wanted to change your introduction, slice the data a little differently, probably run one additional experimental treatment, et cetera, you still will have the opportunity to take advantage of what you learned today and to kind of submit to the special issue.

I should say that we did get 50-plus submissions, as Ginger said. Many of the submissions were very excellent. We could not put all of them, given our nine-paper limit, in the conference, and we were also trying to make sure that we were spreading the topics as widely as we can. So if some of your papers are not accepted, it doesn't mean that it wouldn't do well at the special issue itself. So please continue to work on some of these papers and new papers that

394

might come about. And we're really looking forward to a lot of submissions in the special issue.

About what Ginger mentioned about future things we might be thinking about, something like maybe, you know, outside of Marketing Science of perhaps a biannual conference around this that allows us to kind of bring together people with these kinds of interests, but let us know how you felt about this and, like, whether there is interest. I think it doesn't have
to be me, but, like, you know, as a community, I think it would be great if even after Ginger leaves the position if we can get the FTC to jointly do this and make this an institution and hopefully would be great for the community as a whole.

So thank you very much for the participation and thank you, Ginger and the FTC staff, for your wonderful efforts in putting this together.

(Applause)

DR. JIN: Thank you. In the program, you will see we'll have a dinner at 6:00, which is just a 15 minutes walk from here if you cross the National Mall. The restaurant name is Charlie Palmer Steak, and the address is 101 Constitution Avenue, which is just in the intersection of 1st Street and

395

Constitution and right next to the Capitol. So you should -- you would not miss it. We'll see you there.

If you haven't registered for the dinner, you're welcome to join us. Thank you.

(Applause)

(Whereupon, the conference concluded at 5:37 p.m.)

396

CERTIFICATE OF REPORTER

I, Jennifer Metcalf Razzino, do hereby certify that the foregoing proceedings were recorded by me via digital recording and reduced to typewriting under my supervision; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were transcribed; and further, that I am not a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of the action.

JENNIFER M. RAZZINO, CER

M

9