

[SQUEAKING]

[RUSTLING]

[CLICKING]

**GLENN
ELLISON:**

So today I'm continuing my lecture. I'm starting a new topic on online markets. And in some ways, I think in online markets, it's come to be that Google and Facebook and Amazon tend to really dominate the online space and earn most of the profits and most of the page views or whatever. So I think most of what I'm going to do is try to talk about those leading firms.

Actually, most of today is going to be about Google. I'm going to get a little bit to talking about models, how people think about Amazon at the very end. I actually don't know a lot of good stuff about social media other than there's literature about social media addiction, which I think is also an interesting topic there, but not much theory on it. Oops.

So anyway, before starting today's topic, I'm going to do a brief review of VCG mechanisms. And I think Steven is teaching [14.]124 this year. He'll do a better job. But I wanted to go over this quickly because it helps for thinking about allocating problems. So I'm going to find a private value social choice problem to be given by four primitives.

First, you have a set of types for each player, which you can think of as typically holding their preferences. But it could be other information that the players have. There's some probability distribution p on the Cartesian product of all the θ 's, and we're going to think of those p 's as not necessarily independent across players.

There's a set A of possible social alternatives. In the simple auction example, the social alternatives are give the good to person one, give the good to person two, give the good to person three, and so on. And I'm going to talk about online advertising examples. The social alternatives would be all of the ads that are displayed on the page and where they're displayed on the page. And the complete set of ads are all possible social alternatives.

And this model also does other things like you can have group of n roommates get two tickets to a play and they have to pick two people to go to the play together and the people could care about who they go with and whatever. Then the alternatives would be all of the sets of at most two people who are going to get the two tickets.

And then so the utility functions depend on the chosen social alternative, depend on your own private type, and then on transfers. And I'm going to make them additively separable. So you have this gross utility function for your utility of action a when your type is θ_i , and then there's a transfer. And so it's important that it's a private value problem. You know your own valuation for all the social alternatives.

We define the utilitarian solution to this problem to be the function a_u mapping θ into a that's just defined by a_u of θ is the action that maximizes the sum of the utilities given the true types. So again, like in a second price auction, the outcome that maximizes utility is give the good to the person with the highest value, because you give it to the highest value, that's the maximum of the sum of the utilities you can achieve.

What is a mechanism is going to be a way for trying to achieve an outcome in a social choice problem. A mechanism consists of a set of possible strategy profiles, you know, one set s_i for each player. An action, so basically, you can think of a mechanism as every player chooses some player i chooses an s_i and capital S_i . And then after we observe what actions, what strategies everyone has chosen, there's some function for determining what the action is going to be as a function of those choices. And then there are going to be transfers paid. And again, the transfer that player i gets, which can be positive or negative, is a function of everybody's strategy.

So if we have the social choice problem, we have a mechanism. Then we're going to think about the game where players observe their types, standard Bayesian game, players observe their types. They then choose strategies from the sets s and then the mechanism decides which social allocation is chosen and what transfers everyone gets.

And so the second price auction is an example of a mechanism that implements the utilitarian solution. We run the second price auction game. Everyone bids against each other in the second price auction, and then in that game, I showed it's a dominant strategy for everyone to bid their true value. And then the high value person ends up getting allocated the object. So the second price auction is a mechanism that implements the utilitarian solution in dominant strategies.

The first price auction is also a mechanism that implements the utilitarian solution. There it does it just as a Bayesian Nash equilibrium, not as a dominant strategy Bayesian Nash equilibrium. But the first price auction also implements the social optimum.

Important result due to Vickrey-Clarke-Groves is that we can implement utilitarian action. The second price auction, in some sense, is just a special case of a much more powerful mechanism. And basically, that we can, even in complicated problems like the one I mentioned with the roommates, where everybody cares about who-- they may be jealous of other people going. They may want to go with some people, not want to go with other people. Whatever preferences you give them over those complicated sets, we can still implement the utilitarian solution in dominant strategies.

So how does VCG do this? So I had this function u , which described what was socially optimal given the types θ . I'm going to define this second function u_{-i} to be what's socially optimal if we ignore player i 's preferences. So given a vector of types θ , what maximizes the sum over everybody else's utility? So for instance, in a second price auction, a u_{-i} is to give the object to the person who has the highest value other than player i ? So this is what's the best thing for everyone other than player i .

I then define the externality that player i exerts on others to be what is everybody else's utility when we take player i 's utility into account and choose the optimal action minus what is everybody else's utility if we ignore player i 's utility when we're deciding what to do? So generally, this is going to be a negative number. If we ignore player i 's utility, if we ignore player i 's utility, we're going to do something that's better for all the other players. It could be 0 whether player i is there or not. We give it to player three. But anyway, this is what we find. This is the externality that i exerts on others.

And so then definition of the VCG mechanism. The VCG mechanism is let the set s be capital θ . So every player is just asked to announce their true type. The action is the utilitarian solution, treating the announced types as true. And then the transfers that every player is asked to make is the transfers that are given to every player is the negative externality that they exert on others by having their preferences considered.

So everyone announces their type. We do what's utilitarian given the announced types. And then if you're exerting a negative externality on others by having your preferences taken into account, we make you pay for your-- we ask you to pay your negative externality.

So the second price auction is a special case. In the second price auction, you get the good and then you pay the second highest person's value. That's the negative externality you exert. Because if the winning bidder wasn't there, the second high bidder would have gotten the good and the second high bidder's bid is their utility. Therefore, you're paying the negative externality you're exerting on others, because you're taking it away from the second high bidder. And so that's why you pay the second high bid, not your own bid.

So anyway, theorem VCG mechanism implements the utilitarian solution and is a truth telling dominant strategy Bayesian Nash equilibrium in any private value social choice problem. So important things to know about this, that VCG is just very powerful. You can do this in models with n goods and models where the social alternatives aren't goods at all. The social alternatives are just things that society has to be choosing among. As long as the players know their own valuation, you can get players all the players to reveal that information, at least reveal it enough that you can choose a socially optimal outcome.

VCG is essentially the only way to achieve dominant strategy Bayesian Nash equilibrium implementation of a . In some ways you can just-- I haven't even shown you the proof yet. But in some sense, the converse of the proof is if ever someone wasn't paying the negative externality they'd exert on others, then they could take an action that makes society better or worse off by giving them slightly more or slightly less utility.

VCG mechanism is not budget balanced. Everyone pays this transfer t . The transfers don't add up to 0. If you think about the second price auction, it's important in the second price auction that the money is going to the auctioneer, not to the other bidders. If the money went to the other bidders, then the second high bidder would have an incentive to lie and claim that the thing was worth more to them than it really is so that they got a higher transfer.

And so what happens is the money, in some sense, there's a set of bidders. The money is flowing out of the set of bidders to this outside party, which is the auctioneer. And that's fundamental in VCG that you can't, as I said, essentially the only way to achieve dominant strategy implementation. And it's not budget balanced, so you can't achieve dominant strategy implementation in a budget balanced way.

And then in terms of is this a practical mechanism? The VCG action spaces could be very large. So famous example is that the United States, again, in the early 2000s was auctioning off the right to provide cellular or telephone service in 394 cities simultaneously. And Verizon was bidding and T-Mobile was bidding and AT&T was bidding. And they had two licenses per city.

And so if you think about what's the set of possible social alternatives, it's like Verizon and T-Mobile get Boston, Verizon, and AT&T get New York. Verizon and Sprint get DC, whatever. If you go with every set of those alternatives, 20 bidders, so there may have been 20 choose 2 raised to the 430th power. That's the number of social alternatives. VCG says you submit your preferences over every social alternative and then that's not feasible and then there'll be the second step.

You have to then figure out, OK, if this is everybody's preferences over these 2 to 1,000 social alternatives, then what's actually the best combination? You'd have to then search over that set and find what's the best allocation of the goods to the bidders. It can be completely infeasible if the number of social alternatives is very large.

Questions? So anyway, Steven will do VCG much better in [14.]124. But I recognize many people have not taken [14.]124, so I wanted to do it quickly. Here's a proof. Sometimes you use slides as just lecture notes. If you want to see a proof, there's a proof there.

So now most of what I'm going to talk about today is Google. So Google is the dominant search engine. In some sense, Google has been the dominant search engine for a while. There was a time when many different countries had their own native search engines and those native search engines dominated Google. And so Google was, well, Google came into the United States and there were four or five other search engines that had 100% market share. Google very quickly after its introduction took over the US. And then just one by one by one, it's taken over all other countries.

Seznam in Czechia was like one of the last holdouts, and Google passed them in 2011. Yahoo! was number one in Japan and Taiwan. Naver was number one in Korea. One by one by one, Google picked them off. Google actually passed Yandex in Russia in 2019, but I guess post war, Yandex is back to being the number one search engine in Russia. And then China is the one major country where Google is not dominant, and Google just has a very tiny share in China. But anyway, worldwide market share of Google was reported to be roughly over 90%.

Google is incredibly profitable. Google has \$100 billion in revenues. It has very high profit margins. It therefore makes tens of billions of dollars a year in profits. In some ways, I think Google is even more important than just the fact that it's dominating the search business is that the fact that Google has this enormous revenue coming in, the search business lets it do many, many other things, and in fact, many other online businesses.

Android operating systems, in some sense, Google bought Android and puts out Android phones as the only competitor to the iPhone, in part because it's using the revenues that it gets from searches on Android phones and the profits that it's getting from its other businesses to keep Android as a viable competitor to iPhones.

Other products like Google quickly dominated online email. Google quickly dominated, well, after a number of years came to dominate the browser market. Google is trying many other markets. Chrome OS and Google Docs competing with Microsoft. Google Cloud competing with Amazon Web Services and with Microsoft Azure. Google Meets. I don't know how well Google Meets is competing. But anyway, in some sense, Google has all of this money and Google can plow this money into other markets. And so understanding Google's profits are essential to understanding what's going on there.

The US Department of Justice filed suit in October 2020 alleging that Google violated the Sherman Act in monopolizing the search market. Part of their case is that Google makes very, very large payments to Apple every year to be the default search provider on iPhones. Justice Department is alleging that that's a monopolization, that Google in some sense is actively acting to become a monopoly.

Several states have also sued Google for monopolizing online advertising and for their ad sales practices. Keep being rumors that the Department of Justice is going to file related claims also and hasn't happened yet, as far as I know.

So anyway, some basic facts about Google is every time you do a search on Google, Google runs an auction. That auction determines what the advertisements are going to be. This advertising auction produces very stable results. The advertising is incredibly profitable. And Microsoft has been losing billions of dollars a year running Bing. Despite the fact that Microsoft is channeling \$1 billion a year or more into Bing, Bing has done very little to erode Google's profit margins or its market share.

STUDENT: What does it mean for results to be stable in this case?

GLENN ELLISON: Actually I guess I'll show you what unstable is. So, I guess stable meaning if you do a search on Google, the ads stay in the same place and the firms don't change their bids. So Google runs an auction. The firms seem content to just put in their bids, leave their bids where they are. The ads stay in the same place rather than the firms needing to cycle through the bidding. I'll give an example of unstable in a minute.

So I assume most people have used Google. You're looking for something, a turkey platter. You do a search for "turkey platter." And what you can get is several kinds of sponsored ads. You get an ad up here at the top. Sometimes you see three of those. On shopping links, sometimes you'll get a box like this with pictures. Then you get a bunch of ads over here on the right side for turkey platters. And then down here, you get what people sometimes call the organic or algorithmic search results. These are the ones that advertisers didn't pay for. And this is just Google picking out what's a good turkey platter site to look at.

Maybe you haven't done Bing. But anyway, if you do a search on Bing, Bing is very much like Google. Again, I do a search for "turkey platter." I get some set of ads at the top. I get image ads here. Actually, those are not image ads. So I guess I get ads here, images that are not advertisements, ads here on the right, and then the organic results for telling you you can find a turkey platter on eBay. Start here.

One thing you notice very quickly if you play with Google is that what Google is doing with the ads depends very much on the type of query that you type in. And the less commercial is your query, the fewer ads there are. So for instance, I type in search for a fairly non-commercial term, Math Counts, which is a middle school American math competition. And what you see is just the mathcounts.org site with several sublinks and then Math Counts practice, a Wikipedia article, various other things about Math Counts, but just no ads whatsoever on that page.

And then if I make my search slightly more commercial by putting in Math Counts preparation, then Google decides it's a little more commercial and they put two ads on the screen up here at the top. No ads on the right, but they do have Math Counts preparation at artofproblemsolving.com and amazon.com to get books to prepare for Math Counts. But then otherwise it's still mostly online search, mostly organic.

And then if I picked "shorts" as a word that's more commercial. People type in the word shorts, typically they're trying to buy some shorts. And so this is a very commercial word. So on shorts, you see you get the full treatment of the 11 ads, which is maximum Google was doing at the time I took these. I think may still be true.

So I searched for "shorts." I get three ads at the top, Macy's, Victoria's Secret, Zappos. And then over on the right, I get one, two, three, four, five, I get eight ads over here. Other places you could buy shorts.

And you'll notice that Google has an awful lot of trouble even thinking of what are the organic search results to display. So they do put in there was a movie named *Shorts*. And so they put in a link to that movie. They do decide that maybe Zappos is a place to buy shorts or Zappos, Forever, 21, and Kohl's are the places they think of to buy shorts. I'm doubting those are the three most common places to buy shorts. But anyway, that's what Google picked out.

And then they also have the-- my favorite was the Wikipedia entry in number five, which tells you shorts are "a bifurcated garment worn by men and women over the pelvic area, circling the waist and covering the upper parts of the legs." I find it hard to believe that there's anybody who did the search for shorts because they didn't know what they were. And then that description. I didn't know what a pair of shorts were, but now that I read that, my English is such that that's what they are. But it's commercial, so that what you can see is they've largely turned the page over to the ads rather than the organic results.

And I think what I and others would argue is that the ads have done a better job than Google could do, than the Google's organic search results here. That somehow the ads have figured out that Macy's would be a good place to buy shorts, the Dockers, the Athleta, AE, Loft. In some sense, they've picked out a variety of places where people might actually buy shorts that, in my guess, is better than Kohl's, but I don't know.

Just one more example. This is when things didn't work. So this is a Yahoo! search example from Yahoo! Search from before when Yahoo! was still competing with Google. I did a search for "Paris Hilton." And you look at what is it that when you search for "Paris Hilton" can be. Well, certainly you look at what Yahoo! is able to come up here for ads.

"Paris Hilton tones. Get complimentary Paris Hilton ringtones." Paris Hilton is not a singer. She's not a music-- what is a Paris Hilton ringtone? What makes-- but clearly, they just had an ad for ringtones that they would throw up there almost no matter what you searched for. And they would put Paris Hilton in the name.

And then they do have the "Paris Hilton cheap prices and huge selection." "Paris Hilton on sale." Again, this was an ad that would pop up even if I typed in "shorts." It would say "shorts, cheap price and huge selection." "Shorts on sale." If I typed in "pencils," it would say "pencils on sale." Clearly Yahoo! had not figured it out.

One of the other interesting things about this, though, is that you notice online machine learning always does things that you don't necessarily intend it to do. So what is Yahoo's machine learning thing doing here? It's trying to teach middle school boys to overcome their school's content filters. So clearly it must be that some amount of the searches for Paris Hilton photos are teenage boys trying to find naked pictures of Paris Hilton or whatever.

And what it suggests is that also try "Paris Hilton private parts photos," "hot Paris Hilton photos," because if you type in "Paris Hilton naked photos," your middle school computer blocks that. And so in some sense, Yahoo! is teaching students how to avoid what your middle school computer blocks. Obviously it's doing it because it notes that people type something in, it gets blocked, and they type that in, and then they look at something for a long time. So it's like, aha, that's what I should suggest they search for.

So anyway, Google runs an auction every single search you do. You type in a search, there's an auction occurs. And the auction occurs to determine which ads to show you in between the time you hit your query and the page loads. So that millisecond or whatever it is, an auction has been run.

Here's an example of when Yahoo! was running that crappy engine that I was showing you. This is what bids on Yahoo! for some keyword for click through rates looked like over a one week period. And you can see the bids were doing like this.

What was happening was at the time Yahoo! was running a first price auction for ad slots. They were running a first price auction for ad slots. So on the word "auto insurance," I want to be number one on insurance page. So what happens, Geico puts in a bid and gets to be number one, and then Progressive is like, oh, I want to bid into penny more than that and I'll put in my bid and I'll get to be first. And then Geico bids again, puts a penny above that.

And the firms would just play these cycling games where they would just overbid each other over and over and over again, creating these Edgeworth cycles on steroids. And obviously everybody has to write a bot to update their prices every millisecond or whatever to keep up with each other. And we had this chaotic thing where the firms were spending a lot of time writing pricing bots rather than just figuring out what the slots were worth to them and putting their ads up there. So I said those other things already.

OK, so how the system works is there can be up to 11 ads. The way I'm going to think about this is it's hard for Google to figure out what's-- it's hard for machine learning to figure out which sites have good shorts at good prices. But the advertisers are going to have information about that. The advertisers know whether people who come to their site buy shorts or not.

And so the idea is going to be that this is readily available information for Google that Google can use to figure out which ads to display. Just ask the sellers how much they're willing to pay to have their shorts displayed. That's going to tell you the people who are going to pay more to be displayed are the ones that are actually more likely to sell things to people when they click on the ad. And that's valuable information. That's information that's easier to get out of those bids than it is out of the whole machine learning stuff.

Now, in theory, what we could do is actually figure out the best value page by just having all the bidders say, there are 50 firms out there that would like to display ads on the keyword shorts. You could just ask all 50 of them, OK, think about every way that we could make up this page, which is 50 firms we could put here, 49 here, 48 here, 47, 46, 45. So we just have 50 factorial over 39 factorial and then say, tell me your bid for that composition of the page with you first and Athleta second and Macy's third and so on. And then we would run the VCG mechanism, we determine who's best, and that would be a way to allocate the ads efficiently.

Clearly that doesn't work, because just the bidding space for the firms would be absolutely enormous. How does American Eagle know what their value is for all of these pages? And then how do they send you that message with a trillion different bids? And it's just an infeasible system. But in some sense, theory would say that's what you want to do is just run the VCG mechanism to have all the advertisers tell you what's the best value of every page to them and then you just pick the value maximizing page.

OK, so what has Google done instead? Google, again, early 2000s invented the following. Advertisers submit standing per click bids. So I can enter any search query I want. I could enter "Paris." I could enter "Paris hotels." I could do "Paris" minus "Hilton" to say if it's got Paris Hilton in it, if it's got both words, I don't want it. I can do these bids times persons in the United States. Person is 18 to 39 years old, person is whatever.

But I just have a fairly expressive bidding language where I can say for a person of this type having done this search, this is my value that I'm willing to pay per click to have my ad clicked on. And people just can submit these asynchronously, just any time you just have a list. Large advertisers will have a million standing bids on Google.

When someone types in a search query, Google has this computer science problem of identifying all of the bids that apply to that search query. But anyway, as quickly as they can, they find all the bids they think apply to that query. And then I'm going to say, suppose there are n bids b_1 through b_n that apply to that query. Then what Google does is it multiplies every bid by a quality score w_1, w_2 up to w_n .

And then it ranks the bids from highest to lowest on the product of the quality score and the bid. And then chooses as winners the people whose products are highest and it picks some m between 0 and n . Sometimes it displays 0 ad, sometimes 1, up to n . And then advertisers only pay if you click on their ad. So advertisers only pay if you click on their ad.

If ad k is clicked, then what advertiser k pays is w_k plus 1, b_k plus 1 over w_k , which you can think of as a generalization of the k plus first price auction. So this is the lowest that I could have bid and still beaten bidder k plus 1. So I bid something that beats bidder k plus 1, and they're just saying think of the lowest number you could have bid and still bid in the same position. That's what we're going to charge you on a per click basis.

I'm going to mostly ignore them today. But the weights are incredibly important. I would guess Google may spend \$1 billion a year on their weighting algorithms. Because the whole point is I didn't have the other Yahoo! examples illustrating this. Yahoo! had an unweighted auction, unweighted first price auction in the old days.

So let's suppose you do a search for the word "tennis racquet" on Yahoo! in the unweighted auction days. What's going to come up first on the search screen? You might think, OK, it should be stores who sell tennis racquets. Like maybe Dick's Sporting Goods is number one or somebody who is going to sell tennis racquets. No, that wasn't what it was under Yahoo! It was people who sold ringtones and people who sold pornography.

And you might think, well, why would pornography show up as number one when someone searched for "tennis racquet"? Well, the whole point is it's a pay per click auction. It's not a pay per display auction. And so if someone types in "tennis racquet" and then you put an ad there that says "XXX videos, promotion \$0.99 a month for the first six months free" or whatever, 99.999% of the people who click on that ad are not going to buy the pornography.

But conditional on seeing the ad for pornography and have it be very clear that it's buying pornography and entering your credit card number, if you click on it, then they're going to make a lot of money off you. And they're going to make more money off you than they would make selling tennis racquets, because tennis racquets have a \$20 profit margin. Pornography is 0 marginal cost. And if you get people to sign up for a \$9.95 a month plan, you make \$120 a year off each customer.

And so if you don't have the weights there and you're only paying when you're clicked, the screen fills up with things that have high profit if clicked on and that usually won't be clicked on by accident. And so anyway, the weights there are very, very important.

If you look at the keyword "shorts," I cannot get on the Google screen for the keyword "shorts." Because if I try to bid for "shorts," I can say I'm willing to pay \$10 a click. But they're going to look at my landing page and my landing page doesn't look like a store that sells shorts. They're going to be like, no one is interested in that. We're not going to sell the keyword "shorts" to him unless he has a high quality score.

But from now on, let me just say from here now, I'm going to ignore the quality score for much of this lecture. Because you can ignore the quality score if you're-- once you restrict to Macy's and Athleta and Loft and whomever who are all selling shorts, then they may all have approximately equal quality scores and the quality scores don't matter so much.

OK, so there's a paper. First paper I want to do is Edelman, Ostrovsky, and Schwarz, which developed the following model to think about how this Google advertising system works and why it seems to work so well for Google. So here's their model. So there are n firms. There will be m prizes. So think of these positions on the screen as the prizes. And if you get prize k , you get z_k clicks.

And then I'm going to assume that z_1 is bigger than z_2 is bigger than z_m . So just somehow exogenously, slot one gets the most clicks, slot three gets third, four, five, whatever. So the slots are ordered in terms of number of times people will click on them. And it is true in practice, people click on things on the top much more than they click on things on the bottom.

They're going to assume that bidder i gets paid off z_k times v_i if that bidder gets z_k clicks. So as if each seller has an expected profit they're going to get every time someone clicks on their ad and their total payoff is just number of clicks times payoff per click. And then it's the v_i that's the random variable that's known to bidder i and unknown to Google. And that's why you need the bidding is to get the high v_i people to be displayed.

Suppose that advertising slots are auctioned by the following clock auction procedure. There's a price clock that starts at 0. The price rises until there are only m bidders left bidding to try to be on the screen. And then from that point on, every time a bidder drops out so that there are only k minus 1 bidders left, the bidder who just dropped out gets position k and is assigned a per click payment, which is equal to the drop out point of the bidder in position k plus 1.

So this is how they're going to think about the Google mechanism of submitting sealed bids with the k plus first pricing plan as the outcome of an English auction, just to specify the game a little more. It turns out the game is better specified as an English auction than as a sealed bid auction. But they specify it as an English auction where just prices go up and then every time-- people drop out in order. Once there are only 12 firms left, 11th firm drops out, they get slot 11. 10th firm drops out, they get slot 10. You pay the drop out point of the person below you. So this is just a way of modeling that sealed bid procedure to give us a unique outcome.

Initial observations. This generalized second price auction mechanism that they're describing is not the VCG mechanism. In the VCG mechanism, you're having to give bids over 50 factorial, over 39 factorial possible outcomes of the screen. Here you're just submitting a single one dimensional real number, which is just your pay per click.

So you're submitting your pay per click. That's a very different thing. So the message space is smaller. And message space is smaller and your payment is different. In the Google mechanism, you pay the number of clicks you get times the bid of the person below you.

What's my VCG cost? Let's suppose I take this slot in the VCG mechanism. By taking this slot, I move the person who would have been here down to this slot and the person who would have been here down to that slot, the person who would have been here to that slot, the person who would have been here to that slot, and this person just gets nothing. So by taking this slot, I'm doing a negative externality on the person who would have been here and here and here and here and here.

And so my VCG payment should be the number of clicks that I take from somebody in moving them from spot i down to spot $i + 1$. The number of clicks that I take from someone else to move them from spot $i + 1$ down to spot $i + 2$. I multiply that by their value, because that's how much I hurt them. And then finally, the person who I bump off the screen entirely, they would have gotten z_m clicks and had payment $v_m + 1$. That's how much I hurt them.

So it's not a VCG mechanism. It just has a much smaller bidding space. And it doesn't have the negative externality payments. It has just the bidder below you bid payment, not the weighted average of all the bids of all the people below your payments.

Truth telling is not necessarily going to be an equilibrium when you have more than one bidder or more than one slot available. So I give this example here of suppose there are two slots on the screen. Slot one gives 200 clicks, slot two gives 199 clicks, and there's no third slot. And then imagine you have three bidders have these values. High value bidder has value of 10. Second has a value of four, third has a value of two.

With truth telling, bidder one bids 10. Bidder one gets the top slot. They get 200 clicks and they get a profit of six per click, because they're paying four for each click they get. 200×6 is 1,200. But if instead of bidding 10 I bid three, then I move down to the second slot on the screen. I get 199 clicks instead of 200. But I'm only paying two for them instead of paying four.

So I wouldn't want to bid. Truthfully, here I would drop my bid from 10 and I would rather bid three, because bidding three gets me almost the same number of clicks at a lower price. So this Google mechanism, it's not VCG. Different strategy space, different payments, and it's not VCG like in that bidders-- it's going to be a non truth telling equilibrium in this mechanism. Kind of like the first price auction is a non truth telling mechanism. But bidders are going to have to shade their bids, because bidding truthfully isn't optimal.

OK, so what happens? Much like the first price auction, Edelman, Ostrovsky, and Schwarz show that this model does have an equilibrium in which bidders shade their bids. So if you're not going to get on the screen, you always want to keep bidding higher and higher if you need to, to get on the screen. Because if you don't get on the screen, you get nothing. So all the losing bidders stay in until the clock reaches their true value.

But then once you're guaranteed a slot, when there are k bidders drop out and the $k + 1$ st bid was b_{k+1} , you drop out at your true value minus something. And it's your value minus something that's a fraction of the-- it's something related to the surplus that you're getting and the ratio of the number of clicks of the two higher slots, like the 200 over 199 . 199 over 200 thing. But anyway, it's your bid minus some fraction of the surplus that you're already getting. You drop out early.

And the intuition is you can't bid-- suppose I'm bidding against-- there are two of us. I'm bidding to get on the very top slot on the screen. And there are two of us who are hanging in the bidding trying to get the number one slot, and my value is \$10.

Well, as the clock is approaching \$10, if my opponent hasn't dropped out at \$10 and I haven't dropped out at \$10, the worst thing on Earth for me is if my opponent drops out at \$9.99. Because if he drops out at \$9.99, I'm paying \$9.99 per click, and I'm only getting \$10 a click worth of value. And so I'm getting \$0.01 per click in surplus. And so I'm getting almost nothing.

So if someone else dropped out at six, there are two of us bidding for the top two slots, at some point I want to take my \$4 per click for a lot of clicks in the number two position rather than hang out and risk the other guy dropping out. Any time I'm mad if he drops out, I shouldn't be there. And so there's going to be some drop off point where you decide, like in a first price auction, I want some surplus. So I just drop out and take my surplus in position two rather than risking this guy drops out and leaving almost no surplus in position one. So there's going to be a bid shading equilibrium.

Why is this? If you think about this equilibrium drop out point b^* of v . The first order condition is I'm indifferent between dropping out at b^* of v and dropping out at b^* of v plus Δb . So what's the difference in my payoff if I drop out at this one versus this one if my opponent doesn't drop out in between?

If my opponent doesn't drop out in between, it makes no difference because I'm still just paying the \$6 and getting slot number two. So the only time it makes a difference to me whether I stay in from 857 to 858 is if my opponent is going to drop out between 857 and 858. And so if this is the optimum it must be, I'm indifferent between dropping out at 857 and 858 conditional on my opponent dropping out in between the two of them, which would be that my payoff if I'm the second person to drop out is I get the higher slot and I pay b^* of v .

Essentially he's dropping out essentially where I was going to drop out. Or getting the lower slot where I get fewer clicks, z_k instead of $z_k - 1$, but I pay the \$6 instead of the 857. And so anyway, so it's this local indifference that tells me I'm locally indifferent between dropping out and getting the lower-- at the point where I'm dropping out, I'm locally indifferent between dropping out and getting the lower or the higher slot. And that gives me an expression that I can use to solve for bid shading.

And then sort of a corollary, the bidding formula seemed unexciting. You shade your bid. Remarkable corollary is when players bid using those formulas I just showed you, the Google auction mechanism results in every player making the same payment schedule as they would have made in a VCG mechanism. So it's a non truth telling equilibrium. Everyone shades their bids. But somehow in equilibrium, everyone shades their bids in a way that this exactly reproduces VCG payments.

Why is that? Well, just a simple induction argument. The losing bidders pay nothing. That's what happens in the VCG mechanism. The bidder in the very bottom position on the screen who's bumping someone else off the screen, they pay the number of clicks they get times the bid of the person below them. That's also the VCG externality.

And then just go one level up. What does the person in this slot pay? They pay the number of clicks they get times the bid of this person. So they pay number of clicks in position $m - 1$ times the bid of the person in position m . But what was the bid of the person in position m ? It was v_m minus this bid shading, which was z_m over $z_m - 1$, $v_m - b_m + 1$.

When I multiply these things together, I get z_m minus $1v_m$. And then this cancels and I get minus $z_m v_m$. So the first two terms I get are the number of clicks that I reduce, the number of clicks that I cost the person who ends up in position m multiplied by their value. And then there's a final term, which is this times this times this. I get z_m times v_m plus 1. So I'm getting the negative externality on the person m minus 1 plus the negative externality in person m .

So somehow that indifference condition that determines the equilibrium bids results in the equilibrium payment that I make being the VCG payment. I'm actually paying that sum of negative externalities of all the people up and down the table. And so do I have another slide here, I take it? Yeah, so implications.

The fact that Google's mechanism recreates VCG means that it's an ex post equilibrium. In VCG, you announce your values. It was a dominant strategy. You don't want to change what you announced. In a first price auction, you bid. You see that you've been slightly outbid when you bid v over 2. You want to change your bid. In a VCG mechanism, you make your announcement. You make your payment. You don't want to change your announcement, because your announcement was a dominant strategy.

The fact that this is recreating VCG means you don't want to change your bid when you've seen what the other people bid. You're making the VCG payments given your true value, assuming everyone else announces their true value, everyone plays the equilibrium given their true value. You're making your VCG payment given the announcements that are being made. Therefore, it's an ex post equilibrium. You don't need to change your bid.

And computer science wise, this is an incredibly important aspect to making a system work efficiently, is you don't want everyone to have to have a bot that's repricing their bids every millisecond, seeing what everyone else is doing. By having the VCG payments, it's an ex post equilibrium. You're happy to have your bids stay the same.

Everyone can just put in these equilibrium bids and then let them sit there. And it's not like I'm disappointed that I'm third instead of second. It's like, I put in my bid and I'm happy where I am. If I had ended up third instead of second, that would be good. But given what other people are bidding, I don't want to move.

Another implication of VCG equivalence is that this mechanism is efficient. The winners are the firms that derive the most value from being listed. And this gives you a suggestion for why it may be that the Google ad pages work so well. If it's the case that value of an ad for shorts, the profit that goes to the seller is proportional to the consumer surplus. Because consumer surplus tends to be divided evenly between sellers and buyers in an equilibrium.

Then Google is picking out the highest value advertisers there via this mechanism. Whatever this page max, over all possible 50 factorial over 39 factorial orientations, this is picking out the page that gives the largest sum of values to the firms, which may be the largest sum of consumer surpluses to the consumers.

Third observation is remember I said VCG is not budget balanced. This is what I sometimes call the commercial genius of Google. That is, VCG is not budget balanced. The only way to choose the right ads is to have firms making VCG payments for being on the screen. Where's all that money flowing? It's flowing from the advertisers to Google.

And let's suppose Bing were to say, OK, we would like you to-- we're going to try to steal advertisers away from Google by charging the advertisers less and hoping the advertisers will abandon Google and advertise on Bing instead. This theorem basically says you can't do that. You can't charge the advertisers any less than this and find out which are the best ads to show.

So Google, in some sense, has this immunity to undercutting in price that many firms don't have. The theorem is if you try to undercut Google on price, you're picking the wrong ads and consumers will like the page less. So in some sense, to do this business and to choose the right ads, you have to charge the advertisers this much. And price competition doesn't work the way it would otherwise work.

At one point, Microsoft had, maybe they still have it, they had Bing Rewards. They had a big rewards program for the advertisers where it would be like if you advertised a lot on Bing, we'll give you an advertising credit and give you 50% back at the end of the month for having bought a lot of Bing ads. What's the equivalent of that game? Just everyone bids twice as much, get 50% back, and then they're happy with their position. In some sense, it's hard to undo this logic that if you want to choose the right things, firms have to make the VCG payments. It's hard to mess with that.

So other issues. I think Edelman, Ostrovsky, Schwarz is a beautiful theorem. I think in practice, these weights are incredibly important. The weights are what let you do this pay per click auction. Because without the weights, you couldn't do pay per click and have the right bidders selected, because then there'd be this big bias towards people whose ads don't get clicked on or people don't get clicked on unless for people who-- most people in search are not interested in those ads. And so the weights are incredibly important.

The weights do reflect-- if you look at the weights that Google is using, the weights clearly reflect what people are going to click on. But the weights also reflect much more than what people are going to click on. The weights are also going to have to reflect whether people buy for something after clicking on it rather than they're disappointed after clicking on it and so on.

Google makes liberal use of reservation prices. Many, many pages have no ads whatsoever. That Math Counts page, I've got a book to sell that has middle school math practice for kids. If I'm not willing to bid \$3 or \$2 a click, Google is not going to display it. And they've chosen a very high reservation price for that page. They're not going to show ads that have less than that reservation price.

I think Edelman, Ostrovsky, Schwarz is not an adequate tool for thinking about the welfare implications of that. In Edelman, Ostrovsky, Schwarz, reservation prices are purely consumer welfare reducing. If these firms have values for all of these slots, then by not putting the ad on there, you're just reducing the number of people who click on them, the total. You can't get any consumer surplus if there's no click.

So Edelman, Ostrovsky, Schwarz would say the fact that Google is using reserve prices is evidence that it's exerting market power and market power is always bad. So I think the next paper I'm going to cover is about, in part, how do we think about reserve prices and why they may not be bad.

Anyway, so next paper. Any questions on that before I go on to mine? Athey-Ellison, position auctions with consumer search. The main idea of our paper is to endogenize the prizes in the Edelman Ostrovsky, Schwarz model as the profits that you get from selling to consumer population. And the idea is going to be that we're going to marry Edelman, Ostrovsky, Schwarz with a search model where consumers have search costs that they incur every time they click on an ad and only click on an ad if they think that the ad is going to have something that's worth showing them.

What are motivations for doing this? Obviously, thinking about the consumers clicking seems like a reasonable channel. Think about where the value comes from is going to think about changes you're going to want to make to the standard model. In particular, one idea we have is that reserve prices need not be welfare reducing. In some sense, why are reserve prices not welfare reducing?

Consumers have to always make these decisions. Do I click on ads? If every page you put on-- if you're Yahoo! and every page you put up there, it's like, here are the Paris Hilton ringtones thing, what do consumers learn? Consumers learn that the upper right corner of the screen is where you can find the find x ringtones ad. Consumers won't even read that. They won't read it. They will never click on it. If there actually is something that's useful to them up there, they won't see it. They won't consider clicking on it.

So by using a reserve price, reserve prices can convince consumers that there is probably something. The fact that when you see an ad, you know that person was willing to pay the reserve price, which means they probably have something good to sell you. So you should read it and you should think about clicking on it, depending on your clicking costs. But the idea is that reserve prices can be valuable as a way to signal to consumers that the ads that there are worth clicking on and then can actually increase the total number of clicks rather than always decreasing the total number of clicks.

And then also just we want to think about when you have a model with the consumers, you can't think about consumer and social welfare in Edelman, Ostrovsky, Schwarz because the consumers are just mindlessly clicking and buying things. You don't have the consumer search costs that are being incurred when you put the ads on the screen, what you think about what the consumer welfare is going to be.

So anyway, our model basic elements. Consumers are going to have a need. A need could be a pair of shorts. They can meet this need by purchasing from some advertiser. Consumers incur some kind of search costs every time they click on a link. The advertisers differ in quality. Quality is just going to be the probability of having something the consumer wants to buy. So everyone has some pair of shorts they want to buy. If you click on their ad, they may have the pair of shorts you want, in which case you buy it.

Quality is private information of the advertisers. Some advertisers know that 90% of people who click will find the shorts they want. 30% will find it, 10% will find it. And then if the bids do end up being ordered in a way that's monotone in quality, like the highest quality firm is here, the second highest quality firm is here, the third highest quality firm is here, then that creates guided consumer search where instead of having to search randomly over all the firms, you know to click on the highest quality firm first, then the second highest quality firm, then the third highest quality firm. And that guiding in your search can be, in some sense, it's a intermediary that improves the quality of search by giving you partial information about the qualities.

So anyway, observations in the paper. You can think about these sponsored searches as a means of as an information intermediary providing partial information to consumers, which is this ordering of the firms on quality. We can generalize some of the Edelman, Ostrovsky, Schwarz model to this, even though there are common value elements now. Reserve prices can improve social welfare. And the standard results about click weighting being weighted. Things I'll discuss. You would want to do the weighting differently in this model than you would do weighting in Edelman, Ostrovsky, Schwarz.

I'll do this relatively quickly and skip some of the results. But so the basic model is we have this continuum of consumers of unit mass. Everybody has a need. They get a payoff of 1 if they meet their need. The consumers have search costs s_j from clicking on a link. And then we're going to assume that search costs are heterogeneous in the population. And what consumers do is they search optimally until their need is met or the benefit of further search falls below s_j .

We have n advertisers bidding to be the sponsored links. Firm i has probability q_i of meeting a random consumer's need. q_i is private information of the firms. Firms get a pay off of 1 if they meet a need. And what the search engine is going to do is conduct the standard, unweighted, generalized second price auction, like in Edelman, Ostrovsky, Schwarz. And then they display m sponsored links ordered according to the bids.

So just benchmark for how the improvement in search quality works. What would you do if all you saw is the ads and the ads had no ordering to them, what do you do? You just click on a random ad. You have some expectation of the quality of a random ad. I either don't click on anything or I click one random ad, and then I just keep clicking random ads. Every ad is independent of every other ad, so I just keep clicking random ads to see if they have what I want until I run out of ads.

What do you do if the ads are ordered? If the ads are sorted by quality and equilibrium, then consumers follow a top down strategy. They start at the top and they just continue clicking until either their need is met or until the expected quality of the next website is less than the search cost. And as you move down the screen, expected quality is decreasing for two reasons.

One is, you know that the firms were ordered on quality. So that as you're going further down, you're getting a lower order statistic from the distribution. And then you're also Bayesian updating, because you're learning that if the top three firms didn't have what I'm looking for, it must be that the set of ads is lower quality than I thought it was ex ante, and the realizations were all low. Therefore, the realization of the fifth one is even lower than my ex ante expectation of the realization of the fifth one.

So as you move down the list, your expectation of the probability of meeting your need goes down for two reasons. You know they're a lower order statistic and you know that the top firms were no good, which makes you think that the whole set of draws from the distribution were a low draw. And so as you move down, you're just going to start at the top and then move down until your updated expectation, given that you've had a whole bunch of failures to meet your need is less than your search cost.

What happens in the bidding game? Bidding game is just like Edelman, Ostrovsky, Schwarz. Suppose firms dropped out at these prices. You're deciding how long to stay and accept slot k . Again, you're just going to have this indifference between dropping out at this b^* and dropping out at $b^* + \delta b$.

If you're in the slot $k - 1$, if you're in the higher slot, you get $q_i - b_k$ times the number of slots that go to slot $k - 1$. But what happens if you drop-- I'm bidding against another firm for slot number one. What happens if I take the lower thing? If I take the lower slot, I get $q - b_k + 1$ times my demand for being in slot k . But my demand for being in slot k depends on the quality of all the firms that are in higher slots.

So if I'm bidding against someone and I realize that they have quality one, then I don't want to lose to them. Because if I lose to them, a consumer is never going to buy from me. So in some sense, it's a common value auction instead of a private value auction, because I care about the quality of the firms I'm bidding against. Because if they're really high quality and I get outbid by them, I'm never going to make a sale.

And if they're really low quality, I don't care about being outbid by them, because they're not going to make the sale anyway, so I might as well be in the second position. So if you think about it, if we had the rare draw where all the qualities were almost 0, there's no reason to bid for the second versus the first slot, because the consumers are going to think the second firm is really good. So I might as well just take the second slot with the first guy have 0 quality and outbid me. I'm fine.

It turns out that despite this being a common value auction, the same argument goes through. You can characterize what are the equilibrium bids. It's easy to find the equilibrium bids. So model still has a solution like the other model where there's some bid shading. You shade your bid. You shade your bid because you're going to get fewer clicks for two reasons. Because the posterior is getting worse and because they might have made the sale ahead of you. So you have the two different factors leading you to shade your bid. But it's same deal.

So anyway, but again, I think the heart of this paper, like Edelman, Ostrovsky, Schwarz was the heart of the paper was a theorem about this reproduces VCG. Ours, we have this theorem about welfare effects of reserve prices. And so here's a theorem on reserve prices. Suppose search costs are uniformly distributed. An important property is that consumer surplus and social welfare are maximized at exactly the same reserve price.

And in fact, given any bidding behavior by advertiser and any reserve price policy, the search engine equilibrium behavior by consumers implies that the expected social welfare is just three times the expected surplus or expected advertiser and Google surplus is always twice expected consumer surplus. So there's this social optimum that coincides with the consumer optimum.

And the proof of that is pretty simple. I'm going to find this object called gross consumer surplus, which is consumer surplus plus the search costs incurred. And I'm going to find gross producer surplus to be the advertiser profit plus the fees that Google is charging.

Note that a search produces one unit of each of those if the need is met. Because if the need is met, the consumer surplus is 1 minus the search cost incurred so that the gross consumer surplus is 1. And the gross producer surplus is the advertiser profit, which is 1 minus the search engine fees. So if the search is successful, it produces one unit of each of these. And if it's unsuccessful, it produces 0 units of each of these. If you click and don't find what you want, these are both 0.

So that means the expected gross consumer surplus is the expected gross producer surplus. Welfare is gross consumer surplus plus gross producer surplus minus search costs. So to prove the theorem, we just need to show that the expected search costs are exactly half of the gross consumer surplus.

But that follows immediately from the optimality of consumer search and this assumption that search costs are uniform. Because what's going to happen is every time an ad is displayed, you can think of who's going to click on the ad. The people with a low search cost click on the ad. The people with high search costs don't click on the ad.

What's the cut off? The cut off is whatever information you have, you're going to click if the expected value of the ad, given your information, if your search cost is less than the expected value of clicking, given the information you have, you're going to click. Otherwise, you're not going to click. So what's the average cost incurred by clicking?

It's just the expected clicking cost conditional on your clicking cost being less than this, which is $1/2$ of expected v . So the expected value of search costs, given that search costs are less than a cut off, is just $1/2$ of that cut off. So click by click, whatever all these situations that come up, the search costs are always half of that conditional expectation. So in average, search costs eat up half of the gross consumer surplus.

And so what does that alignment theorem do? That alignment theorem, so we have this. Something that you can think of, if we have of reserve prices here on this axis, there's going to be some amount of consumer surplus that looks like this. And then there's just going to be the producer surplus is just going to be 3 times that. So producer surplus is just always 3 times-- our producer surplus is just always 3 times consumer surplus. There's this unique maximizing point r^* where consumer surplus is maximized and advertiser plus Google profit is also maximized.

Simple example to show that these things can be positive is suppose that there's only one slot on the screen and the distributed search cost is uniform. Then the optimal reserve price is always going to satisfy-- The reserve price is just $1/2$ of the expectation of the quality of the best link, given that the best link has quality at least r . Because it must be that if you're maximizing consumer surplus, it must be ex post.

When you're supposed to display the link, you're supposed to display the link if the bid is above r and not display if the bid is below r . It must be that's ex post optimal for you to do. If I display a link and people think that the quality is bigger than r , this is going to be the-- people are going to click if and only if their value is at least this high. Therefore, the search costs incurred are going to be $1/2$ of this.

And it has to be that I don't want to save the consumers from search costs by not showing it. And so I must be indifferent of showing an ad that has quality exactly r , which means that r satisfies this equation. So if I have one bidder, r is going to be $1/3$ or something like that. I do something kind of like monopoly pricing, and it's because I don't want to-- it has to be that consumers know that this is my reserve price and I only display things that are this good. Then once they know that the quality is high, their expected quality has to be exactly twice r .

Another corollary is that there is this inherent conflict of interest between Google and its advertisers that says that advertiser surplus is lower under the profit maximizing reserve price than under the consumer optimal reserve price. So under the consumer optimal reserve price, we get this consumer surplus. And this is the gross advertiser surplus shared between Google and its advertisers.

Suppose an equilibrium Google doesn't pick this reserve price. Google picks a higher reserve price. So we had this one and this one was shared between Google gets some, the advertisers get some. If Google instead is picking a different reserve price, it must be that Google is getting more profit from this reserve price than it gets at this reserve price. So it must be that the area of Google's profit has gotten larger for Google to want to choose that higher reserve price.

But we know that consumer surplus is lower. Therefore, Google plus producer surplus is lower. So this is smaller than this and Google is getting more, so the advertisers get less. So anything that Google does that's different from the consumer optimal reserve price is something that's hurting advertisers. So there is this-- if your search engine is trying to maximize consumer surplus and then doing something other than maximizing consumer surplus is doing something that's worse for advertisers than maximizing consumer surplus is.

And obviously in the long run, I've talked to other people. We think that it may be reasonable to think that in the early days, Google was trying to maximize consumer surplus. If you're trying to maximize the number of long run customers using your platform, it may be that trying to design the platform to maximize this consumer surplus is a good rule of thumb for doing that.

But to the extent you start monopoly pricing and doing something other than maximizing consumers. So positive reserve prices to maximize consumer surplus, good for everyone. Changing away from that, then its advertisers are suffering. Consumers are suffering. Any departure from this, advertisers and consumers are both worse off if Google is better off.

Anyway, I just think that this is a literature that is still open for people to do other things. Search engine can obviously do better with more general strategies. If you think about the search engine is trying to convey information, you could really bring in all kinds of information design ideas.

You could think about simple things like reserve prices could vary by the position on the screen. The displays could convey more information about quality. I think Google already does this by sometimes putting ads only on the side versus on the top to give you some signal of what the quality is. But you could think about giving more information somehow about sponsored links.

Anyway, the paper also has sort of sections about thinking about weighting. But in some sense, I think weighting- - Edelman, Ostrovsky, Schwarz, every product is independent of every other product. If you think about what Google is trying to do, Google is trying to display a shopping mall. They're trying to create a shopping mall on the page.

So in some sense, what they want to do is also have some kind of diversity weighting scheme, like have something showing women's shorts, some showing men's shorts, some showing athletic shorts, some showing fashion shorts. In some sense, what you want to do is have some weights that reflect how different every ad is from every other ad so that you're in some sense creating the equivalent of a shopping mall rather than maximizing store by store by store what you're trying to do up there.

OK, anyway, I will leave that. People can read in the paper. I have two other things here. I will give you the 30 second version of each of the following two papers.

So one recent paper, Anderson and Renault, "Search Direction, Positive Externalities, and Position Auction Bias." This is addressing a limitation in all of these previous papers, mine included, that in my paper, we take prices as exogenous. Every time you click on an ad, you get one unit of consumer surplus, one unit of producer surplus. We don't let the firms change prices.

You could imagine that in a model like this, firms are going to use their pricing to exploit their position on the screen. And so what Anderson and Renault work out is what would happen in a model where firms' pricing is endogenous? And they can set in some sense, the high price, the first firms, they can exploit their pricing advantage. But they also may want to use low prices to keep consumers from searching the things below them.

And in their model, it turns out that turns out to be the more important factor that I want to set a price low enough so that consumers who like my product don't have an incentive to search the rest of the screen and find the lower prices. I want to take advantage of the fact they've sunk their search cost on me to get them to buy from me and not search everybody else.

So anyway, they have this model where consumers have this need. Every seller meets the need with some probability. But then if consumers meet the need, the gross utility they will provide is something distributed on some interval v lower bar v upper bar. And then the firms set prices optimally knowing their position on the screen. So once they learn their position on the screen, they set their price.

And the interesting aspect of this paper is they note that this model has multiple equilibria. Even if you have just two firms, you can have the low quality firm on top and the high quality firm below or the high quality firm on top and the low quality firm below. And they can both be equilibria with consumers searching top down. Because in each case, all the firms know that the consumers search top down.

Therefore, the top firm sets a really low price to keep consumers from searching the second one, and therefore, having either firm on top at a low price and then the firm below at a high price to only serve consumers who couldn't meet their need at the first firm. Those are both equilibria. And so they work out that this becomes a model where that has multiple equilibria. And different equilibria are better for consumers and different equilibria are better for firms, and they all involve this endogenous pricing. So I think it's an interesting recent working paper.

And then the second one I wanted to mention is Armstrong and Zhou on consumer information in platform design. So this is thinking about Amazon or other platforms who are mediating the process through which outside sellers sell products to consumers. And the thought is going to be that consumers are going to get utility v_{ij} minus p_j if they buy from Firm J, where v_{ij} is a random variable that represents your idiosyncratic preferences for the two firms' products.

And the platform design consists of choosing and committing to some signal structure. How much are you going to tell the consumers about the two products? How much are you going to tell the consumers about two products before they click on the ads? In some sense, by limiting how much information consumers get, you change the price competition between firms. Again, this is an endogenous pricing paper.

And first immediate thought is that the consumer optimal outcome is unattainable. What you would like to do is have all consumers buy the product they like better at $p = c$. But having all consumers get the product they like better at $p = c$ is not feasible, because for consumers to get the product they like better, they have to know which product they like better.

And if you tell consumers which product they like better, then the firms are going to have market power knowing that all the consumers know that they like my product better than the other firm's product and are going to price that into their-- they're going to raise their price to reflect that. So for consumers to get the product they like better, they have to know which one they like better. Then there's differentiation. Then the pricing game results in pricing not at marginal cost. So there's no way to get that fully consumer optimal outcome.

The seller optimal outcome you can sometimes do. The seller optimal outcome is to tell consumers which product they like better and give them 0 other information about how much they like it. Because they have no information about how much better it is, then all consumers are ex ante identical, and the firms can just charge consumers what is your value. If you know there are two products, you know which one you like better, what's your expected value for the one you like better? And then fully extract the consumer surplus. And so that's sometimes possible.

And then they work out, what are you going to want to do if you're trying to maximize consumer surplus versus producer surplus? And what they show is that if you-- in some sense, this is going back to this information design thing. We are just saying what is going to be the posterior? What is going to be consumer's posterior on v_1 minus v_2 given whatever information signal you've given them?

And what you want to do to maximize consumer surplus is have consumers who view the two products differently. I think I like product one slightly more than product two or product two slightly more than product one. And you want to have them have a posterior like this that creates tremendous incentive for the firms to undercut. Because most consumers know that I like firm one only epsilon better than firm two, given the epsilon information you've given me.

And so by giving them this epsilon information about how much they like product A versus product B, then you've created very high elasticities. And so the firms price very low. And even only having this thinking that you like this one epsilon better than the other one in expectation lets you often get the right product. So giving consumers this minimal information about-- lots of consumers, minimal information about which product they like better creates very intense competition and makes that trade off of getting a lot of consumers to buy the product they do like better.

Whereas then the profit maximizing solution, profit maximizing solution is often to-- you give consumers-- many consumers think they like one product much more than the other, and only a few consumers who think that they're roughly indifferent between their products. And you do that to try to get the prices as high as you can possibly get and get many of the consumers to buy the products they like.

But I guess one of the takeaways from this paper was that they found that there's really not much scope to transfer surplus between consumers. Not much scope to modify social welfare. A lot of it is transferring surplus between consumers and producers by giving this kind of information or this kind of information. But it's harder to modify the equilibrium social welfare. Anyway, if you're interested in reading recent papers on Google with pricing or on Amazon, these would be two I'd recommend.

So anyway, Monday I'm going to be back here. I'm going to be doing empirical work on online markets. And these will be actually largely papers on online retail and thinking about consumer preferences and online retail.