[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK. Let me go ahead and get started. So if you remember, last class, I talked about theory of price discrimination under monopoly and did first, second, and third degree discrimination. Today, I'm going to talk about empirical work on price discrimination.

And just as a reminder, as I said previously, all the models I talked about in last class are 100% mathematically correct. There's nothing we need to test about the theory itself. But many things that empirical work can do for us. In some sense, it's asking whether the assumptions are correct and whether things are working like we think or whether there should be other assumptions in the model.

So questions, examples we get help with. Do firms price discriminate? Is the way that firms are pricing consistent with what the models say about how prices would be related to elasticities and other things if firms were profit maximizing? Magnitudes are important. Literature makes a big deal about price discrimination. How important would it be in practice? How large would the profit gains be? Would firms really be incentivized to do it?

I discussed Bergemann, Brooks, Morris, whether consumer surplus goes up or down with price discrimination really depends on what kind of information the firms have. What do those look like in practice? And if you are going to discriminate, what types of information do we need to discriminate?

And I think this is something that's-- as I talked about, the interest in online price discrimination has really increased attention to the field because it used to be discrimination only occurred on expensive items where you could pay somebody to try to discriminate among consumers. Now, in the online world, where you know so much more about people, maybe there's going to be much more scope for discrimination. At least it's an interesting hypothesis.

Early papers on price discrimination. One of the first things that happened in this literature, after the NEIO revolution, was just people arguing, do firms actually price discriminate? How can we tell whether price discrimination is really occurring or if it's just cost-based price differences?

Two of the early papers on this. Severin Borenstein had this very nice paper that just argued-- at the time in the United States, you could either buy leaded or unleaded gas. At some point, we switched from having leaded gas to unleaded gas because leaded gas caused lead poisoning, all kinds of public health problems. Because unleaded gas was initially used by people who had newer cars, those people were wealthier than the people who had cars that ran on leaded gas.

Unleaded gas was a lot more expensive than leaded gas. And so the question was, was that because of the difficulty of preparing unleaded gas, or was that due to price discrimination? Borenstein has this paper, just basically doing very detailed cost accounting and arguing that, no, the price difference between leaded and unleaded gas was much too large to be anything other than price discrimination.

And then Andrea Shepard has this paper, which would be totally self-evident to those of you who are from California. I know in California, you go there, and it's like self-serve gas is $5.99 a gallon and full-serve gas is $7.59 a gallon. And you just do the math in your head. I'm paying someone $1.50 a gallon to pump my gas. That's $2. That person, they only make $15 an hour. There's no way that's what it costs.

But Andrea Shepard has a paper that looks at stations actually in Massachusetts, comparing stations that offer both self and full-service with stations that offer self-service only and full-service only, and using the gap between the self and full-serve only stations to estimate what the cost differences could be. And then looking at the mark-ups between the dual product stations and arguing that that's clear evidence that there's second degree price discrimination going on, where you're charging the wealthier people more for the high-quality product as the gas that someone else pumps for you.

It's instructive. If you want to go back and get a sense for the history of the economics profession and how much the required level of statistics and data has changed, if you compare Shepard's paper in the *JPE* 30 years ago to the most recent paper, last paper I'm going to cover in this lecture, you get the sense of how the economics profession has changed in standards for empirical work.

Second topic I wanted to discuss, though, is racial and gender discrimination in retail prices. This is another place where IO is definitely overlapping with the public policy question. The public policy question is, do prices differ by race and gender? And what's the effect of race and gender-based price discrimination on people in disadvantaged groups or potentially people in advantaged groups?

The early papers in this literature reported very different results. I think one of the-- really one of the first audit studies in economics, Ayres and Siegelman. Ayres and Siegelman, basically the idea was, they just trained 38 undergraduate students who are like juniors or seniors to bargain for car prices. And they train them with a very specific script.

So they were supposed to walk into a car dealer and say, I would like the Toyota Corolla, four cylinder engine, base trim. I'm willing to pay $10,200. I don't need financing. Can I have it at that price? And then they just recorded what did the dealer say back to them. And they took these 38 trained dealers-- of the 38 students they trained, 19 were white males, and the other 19 were not white males.

And they sent pairs of them to the same dealer to ask for the same car, at times far enough apart that it wouldn't seem suspicious. And then they report back, what were the dealers' initial counteroffers, what was the end of the price they got after they had a long bargaining script they were supposed to follow, which included things like, if you say $10,000, he says $12,000, say, how about $11,000? Or if he says $10,000, you say $12,000, how about $10,200 or whatever?

So they gave them specific scripts they were supposed to follow. If you read the paper, they even said, they rated all the students as being of average attractiveness because they were concerned that better looking people will get better deals. So they had white and Black students, men and women. They rated them all as being of average attractiveness in their hiring process, which I don't know whether hiring people by attractiveness is allowed these days.

But anyway, they did that. And that was-- results were that there was tremendous discrimination against women and minorities. So the Black males were the ones who were treated worse. They were charged $1,100 more than white males for a car in the initial offer. And cars at the time were about $10,000. That's like charging 10% more on the price of a car.

Black females were $410 more than white females. And I forget what the number is, but white females were also more than white males. So they reported there was this substantial discrimination, 5% to 10%, in the price of cars.

About the same time, Penny Goldberg had a follow-up article, looking at exactly the same topic, using data from the Consumer Expenditure Survey. And one thing you worry about with the audit study is that these people were not actually buying cars and were the people who were not actually buying cars, would the prices they're bargaining end up being the same for real people who buy cars, and maybe the way that the dealers react to these college students was different from the way they would react to other people.

So Goldberg used car purchases from the Consumer Expenditure Survey and did the same racial test. And she found no difference at all in prices paid for cars. Again, though, data limitation of Goldberg's paper was that, the way the Consumer Expenditure Survey is, is it's an annual survey. And you go to people like me and say, Glenn, did you buy a car in the last year? And I'd say, yes. What kind of car is it? It's a Hyundai Kona SUV. And then what did you pay for it?

And so you're relying on both people describing the car to you and their memory 0 to 12 months later of, what price did you pay for your car? If you asked me what I paid for my last car that I bought, I have no idea within several thousand. At the time, when I was bargaining for the car, I knew very well what I should pay for the car and what a good deal was or whatever. But six months later, do I have any recollection of what the numbers were? I would say I don't. So anyway. But what Goldberg found was that there was no difference across racial groups.

One other study about the same time, Katy Graddy had a paper where she looked at purchases of fish at the New York Fulton Street Fish Market and looked at the same question. There, she did it by just going to the Fulton Street Fish Market, watching people buying fish, and recording what price they agreed to buy the fish at. And she found that actually here, discrimination against whites at the Fulton Street Fish Market, that Asians paid 7% less than whites for the same fish.

But again, these are limited number of cars. Ayres and Siegelman, it's 500 observations. Goldberg, it's not so many. It's more observations, but much more heterogeneity in the cars. And hard to know what the prices are in Graddy's. Again, it's a limited sample. What could one graduate student do hanging out in New York Fish Market for a month?

OK. So first bit I wanted to cover was this one by Scott Morton, Zettelmeyer, Silva-Risso. Multiple motivations. One of them is just to re-examine the earlier findings with much better data. And then besides me wanting to ask the applied question of, Is there racial and gender discrimination in car prices?, it's a little bit of trying to think about, how would we think about this from an IO perspective of IO says that people with different price elasticities are going to pay different prices, and do we see evidence that these price patterns look like they follow price elasticity and other models?

And then the other motivation was, look at this shift from online sale-- from offline sales to online sales and whether online sales are going to change whether we get discrimination, which you could imagine-- you could imagine that in online sales people, the dealers don't know the race and gender of the person buying the car, and therefore they quote different prices. Or you could also get the signaling of things differ-- signaling of knowledge and other things differs. So she's got the cartoon from the paper on the internet. Nobody knows you're a dog. And so dogs can do things different from what they can do in real life.

OK. This is early internet. So this is data from 1999 to 2000. Autobytel was the leading online car purchasing site in 1999. The way Autobytel worked was you would go to Autobytel, you would fill out a form telling it, I want to buy a Toyota Corolla. And then Autobytel would simply send your name and email address to one dealer. And that one dealer was supposed to quote you a good, fair price.

So it was supposed to be a no haggle experience. You just say, I want a Toyota Corolla. They pass your name on to exactly one dealer, and that dealer emails you half an hour later and says, I could sell you that Corolla for $11,300. OK. The nice thing about this paper, though, when you compare it to the earlier ones on discrimination is, the data from autobytel.com, they have the complete Autobytel database.

So 2 million people sent in such requests to Autobytel. They can track whether each of those-- what dealer each of those requests was sent to. And then the more impressive data set, there's a firm called JD Power that gathers data from car dealers. And JD Power is basically-- it's a data provider. They gather information from car dealers about the sales, and then they share that information with other car dealers who want to know what other car dealers are selling things for and with the car manufacturers.

And basically, the JD Power data, it's a complete listing of all cars purchased in the United States from thousands of dealerships and has a tremendous amount of information about those cars. So it's transaction-level data. It has 671,000 car purchases from 3,562 dealerships. On every car, you not just get, it's a Toyota Corolla. You get, it's the Toyota Corolla LX with a 1.4 liter engine.

Here are the 14 accessories on the window sticker that it had installed. It had the cup holder. It had this radio. It had, I don't know, the sunroof. All the accessories that are listed on the window sticker. It has the dealer cost of all of those accessories. So it knows that if you go into the car dealer and the upgraded radio is $399, it knows that radio cost the dealer $127. And so it records all that.

It also records the trade-in in overallowance. People often are selling a car when they're buying a car. And you bring your old car, and the dealer says, I'll give you $3,500 for your 9-year-old car. They actually record the dealer's notation of both how much the dealer paid for the car and what the dealer thought the car was actually worth. So the dealer puts in, I gave a $3,500 trade-in allowance for a car that was worth $3,750 or a car that was worth $2,850. So it's tremendous price data compared to what you got in Goldberg or whatever.

It also contains-- this is not privacy. I'm sure they had to go a while through the IRB for this. The data set has the buyer's home address. It has the buyer's name. And based on having the buyer's name, they make guesses of the buyer's gender and whether the buyer is Latinx or Asian ethnicity. OK.

Anyway. And then, yeah, the transaction data from Autobytel has 2 million requests. Because the Autobytel data also has the buyer's home address and name, they can match it to this data and see whether the person bought from the dealer Autobytel sent them to or whether the person bought a car from a different dealer later.

So initial analysis of this data is very, very simple. It's just a regression. They're just regressing the log of the price paid on demographics of the buyer and a set of controls x. The set of controls has car fixed effects, controls for the cost of all the included accessories. It's got month, region, weekend, end-of-month dummies, a whole bunch of fixed effects. And the idea is that you can just pick up the demographic differences in prices from those demographic coefficients. And most of this is just playing off the JD Power data, not the Autobytel data.

OK. What do they find? Well, they're getting much more precise estimates than previous papers because they have so much more data. And basically what these data say is that there is discrimination in pricing, but that it's a lot smaller than Ayres and Siegelman had suggested discrimination might be in equilibrium.

So examples. They don't know who is Black, but they do know the percentage Black in the census block group where you live. And they know whether the census block group where you live is primarily Black. And they find that this suggests that those in Black areas pay roughly 1.5% more than those living in white areas.

Those in Hispanic areas pay 1.1% more. All these coefficients are circled. I don't know if that really helps. Those in Asian areas are paying 4/10 of a percent less. The Hispanic and Asian effects, you can also see them-- people who live in more heavily Hispanic-- more heavily Asian census blocks pay less. Those who themselves have Asian sounding names also pay less. And women pay more for cars. This one is small. Women pay 2/10 of a percent more for cars on average.

In addition to these demographic differences, you also see differences-- I guess, given the 670,000 car sample, just about everything is significant. And so you get, for instance, those who live in more college educated areas pay less. Those in areas with more home owners pay less. So you get other-- you might have thought the discrimination would go in the sense of, like with the leaded/unleaded gasoline, high income people end up being charged higher prices.

Here, the messages sound like the discrimination, to the extent it's there, looks like it's discrimination in favor of more sophisticated buyers. So the people who have college degrees, own their homes, whatever, they end up paying less for their cars, perhaps because as if they're better at bargaining or have more information or more ability to comparison shop at multiple dealers or are better using the internet, whatever the factor is. But again, these effects tend to be relatively small. They're all in the orders of a fraction of a percent in the price.

OK. Why are these prices differences? You can think of a number of reasons. If you are a Chicago School economist, you can come up with, no, these are all just picking up cost differences because my cost of selling someone a car, there is this cost of, I have a salesperson's time talking to them that I have to recoup, even if the person doesn't buy. And so people less likely to buy might get quoted higher prices.

You also have these effects of, who is going to bring the car into the dealer for service? And perhaps the wealthier people are going to bring the car back to the dealer for service and get charged $199 for an oil change instead of going to the Jiffy Lube and paying $19.99. Therefore, the car dealers are more anxious to sell to them because they know they're going to make the money back later.

Obviously, the price discrimination theory suggests that this could just be profit maximizing against differences, either differences in intensity of preferences or search costs, ability to comparison shop, differences in bargaining skills. Or there could be racial or gender bias. I dislike women. I'm going to charge women more for cars just because I don't like them.

So the data set, unfortunately, while it's very rich, it doesn't have-- there's no exogenous price variation in this data. So there's nothing really to estimate price elasticities here. This is just equilibrium prices paid. So that really knocks out a lot of the opportunities you would like to see of, is the price discrimination model really explaining this? You'd really like to say, do women have different demand elasticities than men? And are those demand elasticities such that you would price 2/10 of a percent different because of them?

So instead, they just have some auxiliary regressions, trying to make the case that they think-- so I guess, I don't know if they're trying to make a case. They come now. Their bottom line is they're going to argue that this looks a lot like profit maximizing, price discrimination, and exploitation of consumers with higher search costs or who are less sophisticated at shopping. So what are some of the things that do this? OK. So doo doo doo doo doo doo doo doo doo doo doo doo.

Yeah. So, for instance, the race and gender effects are only a little smaller in the sample that doesn't require financing. And so they argue that in the sample that comes in and tells you, I have cash, I don't need dealer financing, you're more sure that the sale will go through. The cost base difference of I'm not worried this person isn't going to be able to buy the car anyway isn't there. They notice that the effects are really only slightly different in that effect.

Another one is they look at the subsample of people who are trading in cars. The subsample of people who are trading in cars, if someone is trading in a car, they have a car, and they ought to be just about as good at driving around to other dealers and doing comparison. They have similar opportunities to drive in and do comparison shopping because you know they already own a car.

And what they find is that in this restricting the sample to the people who are trading in cars, the race premiums are about half as big as they are in the full sample. And so they say that that suggests that the differential ability to shop, perhaps because you don't have a car, are less able to shop around, would be explaining at least half of the premium that they're finding to begin with.

And obviously, I think part of-- from this comparison, that the effect is only so big if you don't control for anything, they would argue that, look, if it's down to 0.8 and 0.6 here, if we controlled for additional demographic variables, we would expect the gap to get even smaller. Questions on that?

OK. Then the final part of the paper is they also look at people who bought the cars through Autobytel. And the Autobytel sample, the main sample they look for is people who shop through Autobytel and buy from the dealer that Autobytel assigns them to. All right. I forget whether it was buy through that dealer or buy through-- no. I guess it's eventually bought on autobytel.com the same car that they searched for, not a different.

So you throw out people who searched on Autobytel for a Corolla and then buy a Honda Civic, but they keep everyone who searched on Autobytel for a Corolla and buys a Corolla. What they find is that-- two main results here. Do I have arrows? I do not have arrows.

One is that Autobytel buyers pay less. They pay about 9/10 less than the people searching for a car through traditional means. So this is, again, arguing something about either the Autobytel company or sophistication of people using the internet. The fact that they're using an internet to buy a car and getting a price quote, they end up paying less for their cars.

And then the interactions with the demographics, you'll see that the Black coefficient here was 1.5% but that Autobytel times Black is minus 1.2. Hispanic was 0.7 plus 0.5, is like 1.2. You get a minus 2% for buying it online. So that the-- it seems like the people buying through Autobytel, the race and gender effects have largely been eliminated.

And obviously, in many cases, the dealers are still going-- at least the gender, the Asian/Hispanic things, dealers are going to see the name of the person buying the car. So they would be able to discriminate based on those factors. But they appear not to in the Autobytel data. And they argue that this, again, suggests that a lot of what we're seeing is discrimination against people who are less able to shop or less savvy at shopping.

Any questions on that paper? No? OK. So yeah. I guess one final comment I wanted to make is that if you were to ask, What would be the effect of banning discrimination here, at least banning mean discrimination?, this paper essentially makes it suggest that banning discrimination would not do much at all to dealer profits. Because if you think about it, a feature of monopoly pricing is that d(pi)/d(p) evaluated at pm equals 0. So a change in the price has only a second order effect on the dealer's profits.

And so if you're only changing-- if d(pi)/d(p)-- if you have only a second order effect on profits from a price change and you're banning dealers from having 1% price differences, that seems like that would have a very, very minor effect on dealer profits. Now, this is banning mean price discrimination. It is the case that in the data there are some people who pay 5% or 10% more and some people who pay 5% or 10% less than others. So it could be that the dealers are making a lot of profits from being allowed to discriminate, but being allowed to discriminate in a way that systematically disadvantages demographic groups that can't be important to the dealers.

Questions? OK. So then the second topic I figured I'd focus on with the other two papers was this question of, how large are the gains from price discrimination online going to be? And what would that price discrimination do to consumer surplus and welfare? And again, as I said at the start of this lecture, in the previous time, we know that price discrimination is always good for firms because they always have the option not to do it.

But there was that sort of-- we had the Bergemann-Brooks-Morris triangle that said, monopoly pricing is here. Perfect price discrimination is here. If you allow firms to start price discriminating, we don't know whether you go this way or this way or this way. And this was consumer surplus on this axis, profit on this axis. You let firms start discriminating. You don't know-- you know profits are going up, but you don't know whether consumer surplus is going up, is being flat, or is going down, or is going down by enough to be below the 45 degree line and make social welfare go down.

So in some sense, what these papers are trying to answer is, which direction-- how big are these arrows? Do profits move a lot, or do profits move a little? And in what direction does the consumer surplus change when the profits are changing? OK.

So first, I'm going to cover the paper written by someone who was a graduate student at the time using graduate student accessible data. So Shiller is using information from web browsing activities from ComScore. So ComScore data, it's not free, but it's purchasable. You can access it. He accessed it through Wharton Research Data Services.

ComScore data, it's modified web logs. So he has this. ComScore is a company that, basically, it's an information broker that sells information to firms. What they do is you allow-- ComScore will pay you money. I don't know if it's like they pay you $100. And they will say, can you install this browser add-on? And this browser add-on is going to track everything you do online and send it up to ComScore.

So you have to be fairly confident in the anonymity of that. And obviously, we all know it's very hard to have anonymous-- given that people tend to look at their own page, their own Facebook, or their own whatever, how anonymous is your entire stream of web usage? Can someone figure out who you were and what else you were looking at online?

But anyway, the ComScore panel at the time had 61,000 users who were allowing ComScore to track everything they did and reporting that. And then firms interested in tracking consumers or understanding consumer behavior could buy ComScore data. So Ben bought the ComScore data.

The ComScore data that you get in this public way, it's not the full browsing logs. They realized that was a little sensitive, but what they do provide you with is still a lot, which is timestamped lists of all the domains that you go to. So for instance, this person went to facebook.com. They viewed 83 pages on facebook.com over a period of 57 minutes. Then they went to ebay.com, they did this. Then they went to newyorktimes.com. They viewed 12 pages there, spent 15 minutes. And it just gives you the entire history.

It does also contain-- because it's meant for commercial purposes, it does also contain information on things that they purchased. So if you went to Netflix and signed up for Netflix, that would be recorded in the ComScore data. This person went to Netflix, viewed 12 pages, and purchased a subscription. OK?

The paper focuses on demand for Netflix at the time. Maybe you guys are too young to remember this world. Netflix was not an online service where you clicked on movies and could download them and watch them. Netflix was a service that mailed you DVDs. So you would have a list on Netflix of, these are the movies that I want to watch. Netflix would mail you a movie in an envelope. And then, when you were done watching your movie, you would put the envelope in the mail, you would mail it back to Netflix, and the next movie on your queue would arrive.

So it was kind of like a video rental store, except you didn't have to walk to the video rental store. And you could have different subscriptions like a subscription where you only have one movie at a time or a subscription where you got two movies at a time so that you didn't have to wait. The two-movie subscription, as soon as you watch one of the two, you mail it back. And then the second one can arrive in the mail two days later, giving you time to watch it. You still have the second movie to watch while you're doing that.

Why does he do this? In some ways, this is a very clever design. I think, if you think about somebody who has access to this limited data, and he's trying to understand price discrimination, there are several nice things about Netflix. One is the number of people who purchase something in a year is limited. But he can get everyone who's a Netflix subscriber because if you're a Netflix subscriber, you regularly go to netflix.com and view multiple pages and spend a little time there choosing your movies. And so by seeing whether someone is a regular visitor to netflix.com, you'll know whether they're a Netflix subscriber or not a Netflix subscriber.

You can reasonably treat Netflix as a monopolist. Blockbuster video also had a DVD by mail program, but Netflix had a dominant market share. And you can think of them as monopoly pricing. The fact that 16% of all Americans had Netflix subscriptions at the time means if you're trying to think about estimating demand for something, you need the-- you have 60,000 users, but if only 50 of them have bought something, you really only have 50 observations in some sense.

So the fact that you have 60,000 observations and 16% have Netflix subscriptions mean you've got 10,000 subscribers and 50,000 non-subscribers. And comparing the subscribers and non-subscribers is much more powerful than if you're comparing 60 and 59,940.

And then another aspect is, like I said, people do studies of cardboard boxes. Netflix is a very simple business. You can estimate very well Netflix's cost, which is Netflix has to buy DVDs. They have to have people put them in envelopes, send them in the mail, and then get them back and sort them again. Netflix also had just-- so if you look at Netflix's accounting data and you look at how much they spend buying DVDs, how much do they spend on their mail rooms, you can get very good accounting estimates of, what is Netflix's cost of having a subscriber?

So he treats this as this is a product where we can observe demand well, and we know the cost from accounting data. Big limitation, there is no price variation. Netflix did have some coupons or whatever, but he doesn't know that. He doesn't know what most of his subscribers paid for Netflix because he doesn't observe them buying Netflix. He doesn't know what kind of subscription they have, whether it's a two-disc subscription or a three-disc subscription or something like that. So normally you think of, I need exogenous price variation to estimate demand elasticities.

So I guess the second clever insight that he makes in the paper, besides that Netflix is a good experimental subject, is that you can try to estimate price elasticities here without any price variation. So Tobias is going to talk about demand estimation in a couple of weeks, which can't come too soon, it seems, the papers I'm covering. A basic insight of that literature is, if you have exogenous price variation, you can estimate demand elasticities.

And then once you've estimated the demand elasticities, if you're willing to assume that firms are profit maximizing, then you've got this condition that the firm is maximizing over p, p minus c D of p, or the firm is going to be doing p minus c over p equals 1 over the elasticity of demand.

So if I observe the elasticity of demand and I observe the prices, I can figure out what the firm's costs are without any cost data if I assume that firms are profit maximizing monopolists. What Shiller notes is that if you think about this equation, let's suppose that we know that costs are known and Netflix is profit maximizing. If costs are known and Netflix is profit maximizing, then I know p and I know c and I know p, which means I can infer what the epsilon is.

So having the knowledge of what fraction of people buy lets me get dq / dp because epsilon is dq / dp times p over q or whatever. So if I know p and I know q, then I can also get dq / dp. So that gives me dq / dp. So I can really do what we often want to get out of demand estimation, even without any exogenous price variation, if I have cost variation and profit maximization.

So anyway, what does he do? So again, it looks like a lot of the logit-like demand structures I talked about earlier. He assumes that consumer i's utility from buying Netflix package j is some utility from Netflix minus alpha times the price of the package. So alpha is this price sensitivity. Plus xi j, which is the quality of package j, plus this random error term that makes some people who are similar buy and some people not buy.

And he assumes that this average willingness to pay for Netflix among people with demographics X is Xi times beta. And he's going to think about that Netflix could discriminate by charging different prices to people with different observables Xi. And in his data set, there are all kinds of observables.

So he has 18 demographic variables, 15 variables that relate to the quantity and timing of when you are on the internet, but then he puts in 4,600 other variables, recording things like, how much time do you spend on microsoftnetwork.com? How much time do you spend on the New York Times? How much time do you spend on Gamefly? How much time do you spend on many other websites? And so he basically takes the time spent on the 4,600 most popular websites as a 4,600 dimensional description of you.

And then, in his model, there's only a single alpha. So dq / dp is the same for-- well, not dq / dp, but d (utility) / d(p) is the same for everybody. But because the mean utilities differ, people who have a very high v are going to buy the thing basically almost regardless of what the price is. People with a very low v are highly unlikely to buy it. And so the price sensitivity does differ across people because you have these nonlinearities in how much it's worth to you, and that affects how much you're willing to pay.

And so then he's going to estimate these beta parameters. He's going to estimate the beta parameters via a Lasso-like maximum likelihood estimation. Do people here know what Lasso models are? OK. So Lasso model, you're used to seeing simple regression. How does OLS regression work? I'm just trying to maximize the sum of y minus yi minus beta Xi squared. And this is my OLS estimator of beta. Sorry, I minimize over beta y minus x beta squared. OK?

The lasso model is, suppose you have-- and something you're taught in undergraduate econometrics is if I'm going to do this, I need to have many fewer x's than I have observations. And so I need the-- if the number of variables is N and the number of observations is T, I need N small relative to T. I need T going to infinity relative to the number of right-hand side variables in order to estimate this equation.

What econometricians or computer scientists and econometricians have noticed is that you don't actually need that to estimate a regression if you're willing to make the assumption that most of the betas are 0. And if you're willing to-- the assumption is you have a regression where you have many, many right-hand side variables. Most of the betas are 0. The only thing is, you don't know which ones are 0.

Then what you can do is instead minimize over beta the sum across observations i. yi minus Xi beta squared. And then subtract off the sum over all j of lambda j. I actually want to do this. Minus lambda the sum over j absolute value beta j.

So I'm using i to index people, j to index variables in the equation. So if I instead minimize an objective function like this, where I'm minimizing the sum of squared residuals minus the sum of the absolute values of the coefficients. And because you're minimizing the sum of the absolute value of the coefficients, it's like you're putting a penalty function that looks like this if you use a nonzero beta j.

And so anything where the beta j looks like it has a very, very small effect, you minimize a function like this with a kink by just setting beta j equal to 0. And so what you do in a Lasso model is you estimate a regression like this, and you put in these penalty functions. And you put those penalty functions in in a way that sets most of the coefficients equal to zero. And that is asymptotically a valid estimate under some set of assumptions about how fast the number of nonzero coefficients increases with t. OK?

So he's doing something very much like that in maximum likelihood estimation, where what he's doing is he's maximizing the sum of the likelihood of yi given Xi. And then also subtracting lambda times the sum of the absolute value of the beta j's. So he's doing a maximum likelihood version of that, where you're maximizing the sum of the log likelihoods and subtracting this penalty function for every nonzero coefficient.

And so what his model does in practice, I mean, I'm a little surprised at how many variables he keeps, but roughly he keeps 900 of the right-hand side variables in the equation. But then it's as if the other 3,700 variables get dropped from the equation because the beta estimates are so small that you just set them to 0.

Just to give a sense of what his data looks like, the demographics are things like the age of the people. Is the oldest person in the household, 21 to 24, 25 to 29, 30 to 34? He's got race in there. He's got income in there. And then the browsing habits, he's got things like, what fraction of your internet use is midnight to 6:00 AM? What fraction is 6:00 AM to 9:00 AM? What fraction is 6:00 PM to midnight? How much do you use the internet on Monday, Tuesday, Wednesday, Thursday, Friday, Saturday?

So he has this set of minimal demographics on people, although more than you would often see. At least he does have some set of income brackets. And then he's got all of your 4,600 website visits. And the results come out that the most powerful predictor of whether you're going to subscribe to Netflix is whether you're a frequent visitor to Gamefly.com. I don't know if Gamefly still exists.

Gamefly was a company like Netflix, but that mailed you video game CD-ROMs. So you would just have a CD-ROM subscription to Gamefly. You're playing *Call of Duty 3.* You get tired of it, you mail it back to them, and they send you the next video game on your list. You play that for a month, you send it back, they send you the next one on your list.

But anyway. So subscription to gamefly.com was the most powerful positive predictor. The most powerful negative predictor of having a Netflix subscription was, I don't know, ameblo.jp. Some kind of Japanese blogging site. People, I guess, who apparently visit Japanese blogging sites don't do a lot of watching Netflix movies. I don't know.

People who visit audible.com and rent music also tend to rent DVDs. People who do things like visit jacksonville.com, they don't watch a lot of Netflix. You get the sense it's sort of-- obviously he has the age in there, but there are people who sit around and read local news on the internet and people who watch entertainment.

And it is-- I don't know. I find this interesting, that, yes, once you think about it, yes, many, many, many of these things make sense if I'm trying to predict whether someone is likely to buy Netflix or not. There's a tremendous amount of information in their website. And people are very, very different, and it could be revealed through these many different channels simultaneously.

OK. What does he find? First finding here is that demographic targeting-- and again, remember, he had the same demographics we were talking about in the car paper. He does have-- he has race, he has Hispanic. He's actually got income. All of those things, he says, are basically useless for pricing Netflix.

He said that if you estimate-- I'm going to estimate demand elasticities as a function of the demographic group and think about switching from the optimal uniform price to the optimal demographic targeted price. And he finds that targeting the demographics would let Netflix increase its profits by 0.3%.

I think, though, in some sense, this is consistent with what I said before in that, again, with that sort of-- if profits are quadratic in the extent to which you want to change the price-- and profit, it's a second order change. You only want to make changes on the order of a few percent, then the profit effect from that has to be small. And that's what he says. The profit impact is 0.3%.

He argues that targeting on the web, though, in some sense, may be much more common than targeting in the offline world would be because these browsing variables are much more powerful. He estimates that they would let you increase Netflix's profits by 13%. It's still not a huge number when you think about you've got those 4,600 variables in there. What could you do with all of those? He says it's much, much more than you can do with demographics, but it's 13%.

What are the welfare effects? Oh, do I have it? OK. So welfare effects are-- here he's arguing that with demographic targeting, we're actually-- in this diagram, it's the tiniest arrow in the world, but it's an arrow to the right. That is, the profits are going up by 0.3%. The consumer surplus is going up by 0.05%.

Demographic targeting, it's like this kind of arrow. So it's a bigger arrow, going up 13%. And it's going slightly down on the consumer surplus side. Total social welfare, I assume, has got to be going up because you've got the 13% increase in profit. The only negative, a half percent on consumer surplus. I don't know what the mean consumer surplus is versus profits, but assuming that they're roughly comparable, social welfare has got to be going up in both cases.

He does not have nearly as much data to do this. But he says, if Netflix were to use second degree discrimination instead, discriminating across people by offering the one CD plan, the two DVD plan, and the three DVD plan, there, he argues that second degree discrimination would be much more powerful. And he claims, based on a conversation with a Netflix executive, that they were not doing second degree discrimination at the time.

He said that they were doing-- plans are priced proportional to cost with a single multiplier, whereas secondary discrimination model says you should be charging different markups to different groups. But anyway. He says that secondary discrimination could raise profits by much more and is a much bigger deal than the third degree discrimination they can do.

And here's an example of what he's estimating, the consumer types. If you have no demographics at all, it's as if you're assuming that everyone has the constant mean willingness to pay v for the product. If you put in the demographics, it's basically saying, instead of having a spike here, almost everybody has one of these two willingness to pay that are very close to each other.

And then when you look at the full model with web browsing data, you get this bigger distribution like that. And if you remember back to what I said about Bergemann, Brooks, Morris, it's the shape of this distribution that's going to tend to determine whether consumer surplus is going up or down. And it seems like the shape of this distribution is something to the fact that it keeps them from raising their prices by too much. And that's why that would have to be involved in consumer surplus not going down a lot when you do this.

Questions on Shiller? No? I think I had one more slide. So some comments. Again, I think this is a nice, creative paper if you think about what it takes to write papers as graduate students. He had this sort of-- just uses an off-the-shelf data set, and he has these two different ideas. One is this identification idea that, if I can find a product where I think cost is known, I can estimate demand elasticities without having any price variation in the data. And Netflix is a product that's both very popular and where I can track usage and where I can get the observable cost data.

He also has this other nice observation where I can do this comparison of demographic price discrimination versus web browsing price discrimination. I think these are important applied observations to the extent that, obviously, the limitations of the data may make it-- it's a good data point, but I don't know how much we believe. But to the extent we believe that, it suggests that demographics are of very little value for price discrimination. So maybe we really shouldn't expect to see much demographic price discrimination.

It shows that web browsing can reveal a lot of information. And something else I would say is that it also shows that web browsing-based discrimination is-- one obstacle to discrimination is always people hiding who they are. If you try to have a student-- imagine you try to have a student surcharge at ART. If you go there and you're a graduate student, they charge you 20% more than someone who's not a student. It would be very hard to get someone to prove they're not a student to you. That price discrimination wouldn't work because it's hard to prove that you're not actually in school somewhere.

Here, it would be very hard to distort your web browsing behavior in a way that prevented Netflix from discriminating against you because Netflix can discriminate against you by looking at your visits to 4,000 different websites. And OK, maybe I'm going to stay off gamefly.com in the month before I'm trying to buy Netflix to get my price to go down. But that's just-- they're only putting three times the weight on that as they use on these thousands of other websites.

And so you really would have to do a tremendous amount of distorting your behavior to try to save money on Netflix. And given that it's distorting your behavior on 4,000 other websites to get whatever discount they're giving you, it's just going to be impossible for consumers to hide who they are, other than consumers hiding who they are via some technological means, like some tracking blocker software. But then again, obviously what we don't know is, what's the mean willingness to pay in people who use tracker blocking software? Do they pay even more money?

AUDIENCE: Was Netflix able to see people's browser activity back during this time?

GLENN ELLISON: So yeah. Obviously, there are many data brokers out there who stick cookies on your computer when you visit websites. And Netflix could have made agreements with some of these firms to buy their data and get the tracking data from you. I think in 2006, we know Netflix was-- I mean, they're just basically doing $6.99 if you want one DVD, $8.99. They weren't doing any discrimination whatsoever.

I don't know when those data markets developed, but it certainly wasn't long after this that you could purchase an awful lot of data about users from many-- there are many businesses who have this model of sticking cookies on your machine and then selling the data about you every time you visit their site.

AUDIENCE: Sure.

GLENN ELLISON: Obviously, critical aspects of the paper. It's a very nice observation that you can do this without any price variation, but it takes some strong assumptions to do this. I really would like to have-- he's got-- he had vi minus alpha pj plus xi j plus epsilon ij. He's got the vi is Xi beta. You would really like to have the alpha also be Xi times gamma.

You would really like to be able to get at variation in the alpha separate from the variation in the v's. And maybe we would also like a coefficient here too. Xi times eta times the xi j. Like do people-- also, does their valuation for high quality versus low quality also differ with their demographics? So you've just got limited things that you can do.

And another thing we have is that, what are the properties of the utility function? If you remember, I said, with linear demand curves, we know that price discrimination is always bad for welfare. And so does this model with variation in v's, does it have properties that make it so that the consumer surplus doesn't go down much? Or is that something that he's estimated, or is that something he's assumed?

I don't know so much about the properties of this demand system, but in some sense you'd like to try to get at, what are the properties of the demand system that can do that or not? OK. Questions?

OK. So then the final paper I was going to talk about today, Dubé and Misra, it's an example. This is sort of very high quality data set, very high quality econometrics. Again, the paper you would love to have as a job market paper. Obviously, a big difficulty of this job market paper-- this kind of paper is it shows there is this huge benefit towards data in this world and having connections with firms that can generate such data sets for you and let you be the person who can analyze such data sets.

So the particular application is-- I don't know if people have heard of-- we are back in the modern world now. So this is ZipRecruiter. ZipRecruiter is an employment website. If you're a small firm or a big firm and you want to hire employees, you can sign up for an account with ZipRecruiter, and then ask ZipRecruiter to send you resumes of people who have signed up on ZipRecruiter saying, I would like to be a software engineer. I have these qualifications.

You can have ZipRecruiter send you resumes of people who are currently looking for jobs, who meet the description for your job. The way ZipRecruiter works is it's free to use for job searchers. If you're a job searcher, you can just post your resume on ZipRecruiter.com for free. But then the firms who are looking to hire people have to pay a monthly fee to be on ZipRecruiter.

The paper is dealing with small employers. The small firms are firms that had 50 or fewer employees, reported they were going to search for fewer than one to three employees in the near future. And what ZipRecruiter would do with such new firms when they sign up is charge them a $99 a month subscription fee.

And so you can sign up to ZipRecruiter. You pay $99 a month. And then you post the job that you want to get, and they send you a bunch of resumes that match it. And then you can try to interview or whatever, look over those resumes.

This paper, they had internal access to ZipRecruiter. And actually, ZipRecruiter was willing to run experiments for them. You get the sense that this is, in some sense, a joint research consulting kind of relationship with ZipRecruiter, where they get ZipRecruiter to let them run experiments on ZipRecruiter's pricing, and then they report the results back to ZipRecruiter because ZipRecruiter was interested in the question of, what price should we be charging? Should we be discriminating against firms in charging some firms more than we're charging other firms? And can we make more money that way?

So the way it worked is they had this experiment where if you go to ZipRecruiter on the firm side, you can't just sign up. You have to enter in information about your firm. So what are the things about the firm you have to enter? One-- oh, I guess not here-- is the number of employees that you-- this sort of basic-- to get into the bucket at all, you had to be the under 50 employees and looking to hire one to three people.

And then you have lots of features of the jobs, like where is the job going to be located? What type of company are you? How many people do you need to hire? Does the job provide benefits? Do you need to have a resume? Do you provide mental benefits, dental benefits? What category does the job fall in? Is it an administrative job, a technical job, a programming job, whatever?

It's fairly minimal evidence from a menu. But you've got to do all this before ZipRecruiter will tell you what the price is. And obviously, if you told ZipRecruiter I'm a 50,000-person firm, I'm going to hire hundreds of people a month, they were already discriminating based on that information. And they would say, if you want to hire-- if you're a 10,000-person-- you're MIT, you've got 15,000 employees, you're going to use ZipRecruiter to hire 1,000 people a year, you'll get charged many, many thousands of dollars a month, not $99 a month.

But what they did in the experiment was, they had a month-long experiment where when you signed up, if you answered everything that put you in the low category, the under 50 employee category, they randomized what they told you the monthly price was going to be. And they did it-- $99 was the price they were currently charging. But then, in the randomized sample, some people were charged $19, $39, $59, $79, going all the way up to $399 a month.

So they have this explicitly randomized experiment and an explicitly randomized experiment where they can think about discriminating based on the 133 features they get out of this. What are the 133 features? You get a whole bunch of them out of state dummies. So you get 50 state dummies and some Canadian province dummies. You do get job categories. You get benefits, yes or no. Some of these are binary, some are not. Anyway, they had 133, quote, "features" they could then think about discriminating versus.

So the experiment, it's a big experiment. I think the experiment was everyone who came to ZipRecruiter that month, but that's still a sample of just 7,867 firms. Utility model is very similar to what you saw in the previous paper. It's simpler here because Netflix had many different products that needed xi j's to determine which Netflix DVD level you bought. Here, there's just one. It's just like you go through this process, and it's either 1 or 0. You click, yes, I'll give you my credit card and sign up for $99 a month, or no, I don't want to sign up now.

So the utility is again, Xi beta is your mean utility given your characteristics, Xi, for signing up for ZipRecruiter. They then have the alpha coefficient vary across firms. So you take the Xi multiplied by another coefficient vector gives you your price sensitivity. And then you have an epsilon i again logit error term at the end, where it's just some consumers with the same characteristics are willing to pay more or less to be on ZipRecruiter.

And they estimate this by three different procedures. One is a pure maximum likelihood estimation to estimate the-- again, this is 133 parameters here, 133 parameters here. So 266 parameters to be estimated. 7,867 customers. One thing they do is just pure maximum likelihood estimation.

Again, empirically, people worry about over-- 200 and whatever, 266, is too many parameters to estimate with 7,000 observations. You end up overfitting the data. So they use a Lasso-like MLE procedure again to estimate this, or maybe it's a pure Lasso procedure. I forget. They use a Lasso procedure also to estimate this, setting more coefficients equal to 0.

And then they also do a Weighted Likelihood Bootstrap Lasso, which you can look in the paper. It takes up multiple pages describing what this procedure is. The Weighted Likelihood Bootstrap procedure, the Lasso model essentially sets some number of coefficients equal to 0. And all the inference on it is based on assuming those are all 0. And then what's the uncertainty in estimating the rest of the coefficients?

This Weighted Likelihood Bootstrap model is keeping track of both uncertainty over which coefficients are nonzero and uncertainty over the coefficient values. So it's like giving you a posterior over all the different models that might be true and reflecting that you don't know what the correct model is. You have a posterior over that. And then you also have a posterior over coefficients within the correct model or within each correct and incorrect model.

Typical way these things are done, you have a held-out sample. You estimate the model on 90% of the data, and then see how well it predicts on the other 10% of the data. And they say that from that validation kind of exercise, they find that the Weighted Likelihood Bootstrap estimates are the best estimates relative to the other two. That is, they do a-- the maximum likelihood always fits the data best on the in-sample part. But then when you go to the out-sample data, this one is doing a better job predicting the out-of-sample data than are the MLE estimates. So that's what they recommend.

What happens? So first, this is really a monopoly pricing problem-- a monopoly pricing observation, not a price discriminate observation. But what they find is that ZipRecruiter was not short-run profit maximizing to these consumers. This chart here at the bottom is, how much money would they make? How much money do they make in the first month off the customers who got different prices shown to them?

And you show them $19 a month, still only half of them buy at $19 a month. So you're only getting like $8 a month in revenue. As you go up, $39, $59, $79, $99, you keep making more money because fewer people sign up, but this price is five times that price. It's not one fifth the number of customers. You make more money.

But then you go to the right of where ZipRecruiter was at the time to $159, $199 to $249, and it looks like the profits just keep marching up. So the result is that ZipRecruiter was just charging too little for its service from the short-run profit maximization perspective. And the standard error bars, they're not insignificant on each of these. But clearly, if you're fitting a curve through that, it looks like it goes like this and profits are maximized up here somewhere.

So anyway, they estimate the optimal price would be $327 a month, not $99 a month. OK. Anyway. So that's a test of monopoly pricing, is that ZipRecruiter is not static monopoly pricing before this experiment is run. Incidentally, one of the postscripts of the paper is that seeing these results, ZipRecruiter switched to $249 a month. So then maybe short-run profit maximization model does work once you've hired consultants from Chicago Business School.

So second observation is, there's substantial heterogeneity, both in price sensitivity and in the consumer surplus. So as I said, they've got heterogeneity both in the v's and in the alphas. And the paper says there's a lot of both. So this is the fitted price coefficient, what I was calling the Xi times theta alpha, the predicted price coefficient here.

And you can see, on an absolute scale, this is 0.02. So this is like 0. So this thing is double the price sensitivity of this. There's just a great deal of heterogeneity in that price coefficient. So there is substantial number of people who are-- it's around 0.3. There's a lot who it's around 0.6. There's a lot of difference in what that price coefficient seems to be. And again, this exogenous variation in price means we can estimate these things pretty well.

What is the heterogeneity in how much consumer surplus those who are purchasing are getting? Again, their estimate is that there's this big upper tail of values and that some people are getting very little surplus and are on the margin to buying and not, but most people are getting at least $50 worth of surplus. And then there's this thick fat tail of people getting very high consumer surplus from this product.

So they're estimating that there's a tremendous amount of consumer surplus, where consumer means firm. But firm surplus from using ZipRecruiter at the time. How much surplus is there? Well, they estimate that if you look at what optimal personalized prices would be, the optimal personalized prices would range from $126 to $6,292 a month. The $126 is noteworthy because they were only charging $99. So this would say, even if someone has all of the attributes that make you think they're the most price-sensitive consumer with the lowest value, what would you charge? The answer is $126.

$6,292, it's a little heroic there because they didn't experiment with any price above $399. So having an experiment where the maximum price is $399 and then you say, OK, given that experiment, $6,000 is what you should charge that guy, obviously you've got all kinds of functional form assumptions for what happens out of sample, and is that price really correct?

So I think, in fact, when they see things-- one thing they do say is, the mean price would only be $277, whereas its $327 with a uniform price. So the skew is such that you do charge a lot of people low prices. And even with the $6,000s mixed in, you do a lower mean price with discriminating than you do without discriminating.

And the I think more reasonable treatment of this data is, let's look at-- suppose we look at the personalized prices and just cap them all at $499. And when they actually-- ZipRecruiter said that to them, is that we're not going above $499 no matter what you say. So they have this counterfactual of, let's suppose they discriminated, but they let prices go up to $499 based on this prediction.

The model estimates were that discriminating would increase profits by 8.2% relative to the $327 price. So if you discriminate, that's the magnitude of discrimination. So obviously it's much bigger than Shiller's number. But I think of it as still relatively small.

What would the optimal personalized prices do? The optimal personalized prices are the ones that include the $6,000 numbers in them. They say the optimal personalized prices would reduce consumer surplus by about 25% and would reduce social welfare relative to the uniform pricing.

Why is that? Well, one thing. The summary, essentially, of the report is, under these discriminatory prices, even though the mean is $277 instead of $327, demand goes down. And the reason is, they're taking all of these super high value people and monopoly pricing against them. And as I said, when you do monopoly pricing, no matter how high demand is, monopoly pricing, you're always, in some sense, selling to half the consumers.

And so if there are all these high value people out there willing to pay thousands of dollars a month and you still only sell to half of them, there's a huge amount of deadweight loss that comes from not selling to-- taking every one of those really high value populations and not selling to a lot of people in those. So anyway, the estimate is that consumer surplus goes down by 25% here. And so total social welfare has gone down.

Again, profits only go up by 8%. Consumer surplus goes down by 25%. And then how big was consumer surplus relative to profits? Well, it looks like consumer surplus was actually quite big relative to profits. OK.

It's a nice paper in that they actually include my Bergemann-Brooks-Morris triangle themselves. So I don't have to draw in the arrow. So here's what they say personalization does. So personalization takes you from here, and it's headed in that direction. So it's going to the direction where you're really getting only-- if this is the gap in there-- they've done this for their data. So this is the gap between uninformed monopoly prices and what they think would happen with first degree price discrimination.

And so first degree price discrimination, you'd make a huge increase in profits. You're only getting a little bit off the floor by doing this personalization with 133 variables. And you're getting it in this arrow that's going very much horizontally, is their claim, that consumer surplus is just going down rapidly, and profits is only going up a little in this application.

So in some sense, this is one where this is-- it's a positive and negative message from social perspective. One is that it's saying socially, price discrimination is really bad in this application. It's destroying an awful lot of social surplus just to raise a tiny bit of profit. But then the positive counterpoint to that is that given the profit gain is so little, maybe the firm wouldn't even bother to do it because it's only an 8% increase in profits.

And what are the other downsides of doing it? Are you going to pay the consultants to do this and take whatever political pressure, bad news you might get, or whatever to do it? And in fact, follow-up experiment is, ZipRecruiter, in fact, seeing these results, did not do it. ZipRecruiter said, thank you very much. Good to know our prices are too low. We'll start charging $249.

And actually, a very nice part of the paper is that ZipRecruiter then let them come in and do a follow-on study, judging, you have all these very out-of-sample predictions of what things would do. Let's see you discriminate in practice. Again, a held-out sample. So they use experiment number one to predict-- to find the beta hat or theta alpha hat.

And then that lets them get, what is pi hat for every consumer i? Let's see how well does your price discrimination scheme work, reflecting the fact that there is estimation error in these things and you're trying it some number of months later. Does your model still tell us who's price sensitive or whatever?

So anyway, they ran a second experiment. In the second experiment, $249 had been the standard price before the experiment was run. They experimented with 25% of the subjects going back to the $99 price, 25% keeping the $249 price, and then 50% of the sample being on the personalized prices that varied from $129 up to $499. They rounded everything to the nearest 9.

And what you find is that, really, it's impressive in showing that their model is pretty accurately predicting what would happen under these three policies. So what their prediction was was that the mean probability of signing up would be 26% in the $99 group, 15% in the $249 group, and only 14% in the personalized group where they're charging high prices to some of these high value populations.

And what they find in the data, just a simple comparison of means, it was actually only 23% instead of 26% in the $99 group, it matched 15% in the $249 group, and it was 15% versus 14% in the personalized group. If you look at the revenue predictions, they were predicting the revenues, per consumer, would be $25, $38, $41. It's $22, $38, $41.50.

So I think this experiment gives you a lot of confidence that, at least in the short run, yes, price discrimination would raise ZipRecruiter's profits by 8% using this feasible price discrimination scheme. And obviously, ZipRecruiter, I believe, having seen the results of this, still just stuck with the $249 price and didn't go ahead and change and adopt this change to try to get that 8% out of people.

OK, some comments. Again, benefits of working collaboratively with firms is just huge here. It's very nice that Shiller could be clever and think about his identification ability and the ability to do this without any price data, but price variation is good. So obviously, can you walk into ZipRecruiter and get them to run this experiment for you? Probably not.

But if you have friends who have internet startups that are still small enough they would rather hire a graduate student than some Chicago Business School marketing professors to do their experiments for them, finding those opportunities to get data sets like this and let someone run experiments. And often, actually, you can go to MIT alums or people who work at not quite so fancy companies may help out. Wayfair is in Boston. Maybe you have a friend who works at Wayfair. They can get you an introduction to their data science team. Could run experiments through firms like that.

I have to say that segmentation analysis here, that second experiment really shows you that the Weighted Likelihood Bootstrap Lasso thing that they're doing seems like that is doing a good job of segmenting the consumers and predicting their demands. And at least they don't give you so much sense of like, are the price elasticities estimated from the new data the same as they would be from the old data? But at least it looks like this is a very good method for segmenting people.

And if you wanted to think about doing another study, I think of the applied results as being similar to Shiller's in that demographics are just-- they have more demographics than Shiller did. They're using 133 features, whereas he was using 33 features. They're getting more powerful results. Maybe there's a sense of having done the demand estimation better. But it's still not all that big.

Another concern, as I said, going back to Shiller's paper, is it's a very nice study, but it is-- got this maintained assumptions that people have these logit demands with the logit demand errors, and it's only heterogeneity in this coefficient and this coefficient. Again, we would like to know more about, to what extent are these predictions of the model versus results estimated from the model?

I mean, I think it's nice as an empirical result that Shiller was finding this and Dubé and Misra find this. It makes you think that maybe this model is fairly flexible in which direction those arrows could end up pointing. But it would be nice to know more about that.

One other thing that's important-- maybe this is more of a practical consulting point than a paper. But these experiments only identify short-run individual effects. They have a bit in the paper where they look at, do people persist on ZipRecruiter for more than just the sign-up month if you charge them $249 instead of $99? But there are all kinds of things that these estimates are never going to-- these sort of what we call A/B testing in models like this are going to not tell you, which is, what are the general equilibrium or systemic effects?

Is there a word of mouth customer acquisition effect, where someone signs up for ZipRecruiter at $99, they're satisfied with ZipRecruiter, they recommend ZipRecruiter to someone else. Whereas the $249 people are signing up for ZipRecruiter, but they're not telling their friends to sign up for it. Does the price per quality affect when the experts are writing articles in the newspaper about which employment site should your firm work with. Is there some possibility that there will be popular outrage?

Also, we've dropped the number of firms here. It was 23% going down to 15% using the site. In some sense, that means there are one third fewer jobs on the site from small new employers. One third fewer jobs is going to mean fewer resumes posted to the site. Fewer resumes posted to the site, is the thing less valuable to employers? Is there a spiral down where fewer people sign up for it, therefore fewer people use the site?

Thinking about it. Also, I think John is going to do in a few weeks two-sided market models. But this is a two-sided market model where you shouldn't be doing the one-sided monopoly pricing. You should think about, we're setting our price. It's a multi-product thing where I might price to firms and price to consumers. My price to consumers is fixed at zero, but I still want more of them to be there to make it more valuable for the firms. I should be thinking that multi-side thing into effect.

And actually, final thing is, this is a feature from several years ago where Jeff Bezos, back in the early days of Amazon, Amazon appeared to be price discriminating against consumers and charging different people different amounts, just like there's outrage against price gouging for gasoline during bad weather. Anyway, Amazon get into a tremendous bad press for discriminating against people and charging some people more than others.

And Amazon had to come out and say, OK, we promise we're never going to do this. Again, we don't know, are there going to be similar effects-- is this part of why ZipRecruiter said 8% is not enough, because they don't want to get caught price discriminating against people?

OK. That's what I've got for today. So Wednesday, I'm back to oligopoly competition. It's also going to be a very-- back to a mostly textbook chapter where a lot of it is old. I'll get to some new stuff at the end. But I think actually the new theory gets pushed into next Monday's lecture because I'm doing less empirical and more theory on oligopoly pricing.