

Seventeenth Annual Microeconomics Conference | FTC | November 14, 2024

Viola Chen:

Good morning. Good morning. My name's Viola Chen and I'm one of the staff economists here at the FTC. I am also one of the conference organizers along with Sam Kleiner. Welcome everyone to the 17th Annual FTC Microeconomics Conference.

So how many of you took an airplane to get here? Quite a few of you. And so you know at the beginning, before you go on your flight, they have this safety briefing for everyone that you're supposed to pay attention to, but you really don't. I have the pleasure of doing a similar sort of thing, but I do hope that you pay attention, because there will be a quiz at the end and it will count for 50% of your final grade.

One, please silence your mobile devices and any other electronic devices. Two, if you leave this building for any reason, you will have to go back through security. It does not take as long as the airport security, but it will take longer than you think it does. If there is an emergency that requires us to remain in the building, there will be instructions provided over the PA system. If there's an emergency that requires us to evacuate the building, we'll need to exit through the Seventh Street exit, turn left, cross East Street to the FTC assembly area. That's where the church is, and we'll remain there until we have clearance to return to the building again.

For visitors, when you came through security, you got a plastic FTC visitors badge. We reuse these badges for multiple events, so if you could please return those to the security desk at the end of the day, that'd be greatly appreciated. And also for our name tags, which you all experienced the wonderful snafu we had in the morning, if you could return the plastic badges to the registration desk, that'd be great as well.

If you notice any suspicious activity, please alert building security. Please be advised that this event is being photographed, recorded, webcast. By participating in this event, you are agreeing that your image and anything that you say may be posted indefinitely on the FTC website and on any one of the Commission social media sites. The resto phed, rend ony dnd on thathat o phide, m (d in)2ded, wen5-3.996 (t)

-

Okay, so got through that and now we can actually begin our conference. I have the honor of introducing our Bureau Director, Aviv Nevo. We are very grateful for his leadership here at the FTC as he's currently on leave from the University of Pennsylvania where he is the George A. Weiss and Lydia Bravo Weiss PIK professor with appointments in the Wharton School and Economics Department. I give you Aviv.

Aviv Nevo:

Thank you, Viola, for doing a great job with the safety instruction and for getting through my title, which I never can. So, thank you. So as Viola said, my name is Aviv Nevo. I'm the Director of the Bureau of Economics here at the Federal Trade Commission. I would like to welcome you all to the 17th Annual Microeconomic Conference hosted by the FTC. Personally, it's a great pleasure for me to be here.

I was actually on a panel on the very first conference in 2008 that was actually held at the New Jersey Avenue location. I think there's some folks here who might remember that location, but most of you probably don't even know where it was, and I was on the Scientific Committee of the next three conferences after that. So it's really a true pleasure to see this conference be as successful as it is and really sort of continue. Actually, as a side note, I think we've liked this conference so much that we started another one, more of a marketing conference. We had the second annual one just a few weeks ago, and we're going to continue on having it kind of on a spring schedule from year to year.

From those of you from outside the FTC, I want to say a few words about our agency and the Bureau of Economics or BE as we like to refer to it. As you probably know, the FTC is an independent agency and it has two primary enforcement missions. Consumer protection and competition. BE supports these two missions. We have just under 120 FTEs with roughly 95 of them PhD economists, many of them which are here in the audience and you'll get to meet today. That makes us one of the larger groups of microeconomic economists in the federal government and we do a lot for the agency. We support the competition and consumer enforcement missions. We provide economic analysis in support of investigation and litigation, and we apply in many cases cutting economic analysis, both theoretical and empirical.

I've been very fortunate to have the opportunity to work with the BE folks and the FTC staff more broadly over the past two years. At some point, hopefully many years from now, when I look back at my career, serving as a BE director will surely be a highlight. For the young and maybe not so young folks in the audience, if an opportunity to serve at the FTC or for that matter, our sister agency, the DOJ, ever presents itself, my advice is to take it. If someone offers a job, take it. It's a great job. This is truly a unique experience and one that I highly recommend to anyone to do if they have the chance.

The last few years have been a particularly exciting time at the FTC. We've had a lot of interesting things going on and BE is right in the middle of the action. One of the many ways that BE contributes to the agency is by bringing knowledge from the academic community into our work. Interactions with the academic community, like at this conference that combines cutting academic research and discussions of real world policy problem is key to achieving this goal. You might not realize it, but your research can have a real impact on both policy and litigation outcomes. I would like to take a few minutes about a couple of issues that stand in the way of research achieving its full potential impact in this way. Let me stress that the views I'm about to express are my own and do not necessarily reflect the views of the Commission, any individual commissioner or that of the FTC staff.

So of the two issue I want to raise, the first one is we, the FTC and BE, value an open exchange of ideas between the staff and experts outside the agency. This enables the staff to learn about new ideas and perspectives, hone the professional skills and receive feedback on their work. However, such a dialogue requires transparency with respect to relationship with interested parties so that there's no question

whether undisclosed relationships may have extracted influence on the opinions expressed. The profession has come a long way with disclosure requirements set, for example, by either the AEA or the NBR, but we still have a ways to go. I've been to many conferences where presenters do not offer disclosures or make some vague statements such as reference to a webpage. Okay, you want to see my disclosure, it's on my webpage that you can go look at later, but the audience can't see in real time and can't make judgment in real time about the opinions being expressed.

Now in part, I think this is due to lack of clear standards. At times, folks do not know what is expected to be disclosed in oral presentations. For this reason, I would like to share with you today that BE has a new disclosure policy that we're about to introduce and that we will enforce from now on in our conferences and seminars. And frankly, I hope that the rest of the profession follows it or some version very close to it as well. The policy is a modified version of the AEA Journal Disclosure Policy. For example, we clarify the center funding, not just personal funding needs to be disclosed as well as

screen the great submissions and work with the Scientific Committee. So with that, I'd like to turn it over to Sam, who will introduce our first presenter. Thank you.

Sam Kleiner:

All right, thank you very much, Aviv. So to kick off this first session, we have Yanyou Chen from the University of Toronto presenting, Driving the Drivers: Algorithmic Assignment in Ride. And that paper will be discussed by Nick Buchholz from Princeton University.

Yanyou Chen:

Okay, should I get started? Okay, cool. Thank you very much for having me. So I'm Yanyou Chen from University of Toronto, and this is a joint project with my colleague Yao and University of Toronto and Zhe Yuan from Zhejiang University. Following Aviv's instructions on disclosure, this project is funded by the SSHRC in Canada, so Social Science and Humanities Research grant, and I have nothing else to disclose. Okay, so I will present Driving the Drivers: Algorithmic Assignment in Ride.

Recent years, we have witnessed this advancement in algorithmic technologies. While the existing literature has focused mostly on pricing algorithms, especially how algorithmic pricing is going to lead to market outcomes such as collusion, there has been less attention on pricing algorithms, which the platforms also use to, say, shape the worker behavior. One prominent example of those pricing algorithms is actually assignment algorithms, which are widely used by gig platforms, including ride hailing, food delivery, and parcel delivery services. How those algorithms work is usually the platform will assign a score to a particular worker based on the worker's historical performance, and the workers with higher score will receive better or more other assignments from the algorithm. While those systems are designed to optimize efficiency, but they can on the other hand also limit the flexibility and autonomy of the workers. This brings us to our project. So we provide the first empirical study of such a preferential assignment algorithm and we examine what are the impact of those algorithms on like labor behavior and their welfare. Our research centers around two research questions. First, we want to

extract the entire driver surplus, so the preferential assignment algorithm is valuable for them. And another reason is that the prices and wages are nowadays very sensitive topics and they will receive scrutiny from federal agencies and also workers. In a recent case in Canada, the Uber drivers, they complained and had a protest over the new pricing scheme of Uber. And also, this is confirmed by our interviews with the drivers.

So based on our interview with them, we asked them what are their thoughts on, say, search pricing and the preferential assignment algorithm. Their perception is any types of wage differential, they think is unfair. So they should get paid exactly the same wage. But then we asked them, "Hey, what do you think about preferential assignment algorithm?" They think it's fair. The reason is because they think the more you work, then of course you should get higher priority. This boils down to our key insight of this particular paper. So those preferential assignment algorithms are widely used by the platform, and on the surface they may seem fair, but there comes with hidden costs and I will show in more details that they will limit the labor flexibility and autonomy and those often goes unnoticed by the workers themselves.

So given that, let's move to our empirical study to show exactly what is the welfare impacts of those preferential assignment algorithms. Our context is a very large ride-hailing platform in Asia, and we have all the completed and initiated the proposed but unmatched transactions in a given month of this major city in Asia. So we observe every attribute for the order, like the departure, destination, distance, price, so everything you can imagine for a particular order. On top of that, we also obtained driver attributes, demographics such as age, gender, and the birth location of the drivers. For our focal platform, in this particular market, the focal platform has over 90% of market share. So in our study we treat it as essentially a monopoly in this particular market. We understand the preferential assignment algorithm may have some competition implications, like you want the drivers to be loyal to your platform, but this doesn't directly apply to our case. So we abstract away from this competition effects.

Regarding the summary stats for the driver hour level, we aggregate everything to the hourly level so we can construct the hourly wage for the drivers. One thing I want to emphasize here is among each working hour, the driver only spends 30 minutes to serve the drivers, which is also confirmed by a study using Uber data. You can see the idle time is about 15 to 20 minutes. So it means this assignment is really important for the drivers to earn higher profit. First thing we want to examine is who earn higher hourly wages in this case. The platform tells the drivers how the score system works, but does this really result in any difference in the wage? So we want to confirm that. We show that if the drivers work more hours in a given month and their hourly wage or their hourly earnings will indeed be higher, especially if their percentage of working hours is during those incentivized hours. You can see the hourly wage, like earning is higher for them. So we control it for like this. So there is no search pricing for this particular platform. Therefore, you shouldn't worry about, hey, during incentivized hours, maybe per order you get paid more. It's not the case.

And here is how we think about the decisions of those drivers. During our studied time period, the incentivized hours are midday from 10 A.M. to 4 P.M. and from the evening hours from 7 P.M. all the way to next day. We think the decision of drivers and choosing either to be essentially ~~the~~ full worker or parttime worker. So if choosing to be ~~full~~ time, they can commit to one of the 16 schedule, which means they work for at least two consecutive hours during the incentivized hours. That's a minimum requirement. So they fulfill that and then they can freely choose for other hours of working whether they want to work or not. And for the, we call it ~~low~~ score drivers or ~~no~~ committed drivers, they have full flexibility of choosing whether to work or not for each hour of the day.

And then we show some summary stats for who are those ~~high~~ drivers and who are those low score drivers. We found 70% of the ~~locals~~ in this particular city, they have this residence permit,

which you require to get some health benefit to buy houses. So we find for those ~~tools~~, they are more likely to be high

Indeed, the high score drivers drive a little bit faster, but this only explains for 0.5% of this wage differential. So the large amount is still explained by the more number of orders the high score drivers get assigned.

So from here, we hope that we rule out some of the concerns for the drivers. They strategically choose where to work, strategically cancel orders, and simply drive faster. Okay.

After establishing that, we want to know, back to our original question, what is the effect of this preferential assignment algorithm? So if we eliminate this, who will benefit and how large will be the labor welfare become? In order to do that, we build a model of this dynamic labor supply decisions. So drivers will choose whether they want to be high score drivers or low score drivers. And then they decide for each hour of the day whether they want to work like each hour of the day, whether they work or not. Okay.

So for the high and low score drivers, their wage composition comes from first the platform charges 20% of the commission. So it's fixed. The platform doesn't do any wage discrimination. This is also confirmed by our interview because the drivers really hate that. They don't want any wage discrepancy.

So the wage depends on the commission rate, and then this assignment, ST, depending on the platform, decides how much orders they are going to assign to high score drivers compared to low score drivers. Okay. And also this will be effect by this congestion effect like how many drivers are available in each particular hour of the day.

So the workers' schedule for the high school drivers, they can commit to one of those 16 schedules I showed earlier. So they first commit to this schedule and then condition on that for each hour of the day, they decide whether they want to work or not. And for those low score ~~original~~ workers, they can choose freely what they want to do. And then given these choice decisions, then we can construct

So here, you can already see this effect of the high warming up cost. So for the red line, it's the high score drivers. For the blue line, those are the low score drivers. And the solid line means they have worked in the past hour. Okay. So conditional, and they have worked in the last hour, their probability of continuous work is much higher than if they didn't work in the last hour.

Here is our model feed for all those four results. So given the number of parameters we have, we will say we did a good job in feeding the data. The place where we are mismatched is during the early morning hours, like from the 12 AM all the way to six AM. One reason is because you have much fewer transactions during that time period. So that gives us less fitness in that time. Okay.

Our first results for the reservation values, this is the average reservation value for the 24 hours of the day. You can notice during the morning and afternoon peak hours, actually, the drivers, they have lower reservation values, which is very intuitive. But during the midday and especially during the early morning, they have high reservation values. This is one of the reason to explain why the platform wants

So here's a brief or radically condensed summary of the model that we can use to motivate this. Consumers have CES demand. They see the price in the market, and they come to the market. Consumers are sort of second order in the model. It's really about understanding the labor market side. The platform then has a couple of choices to make. It chooses what price to set in the market. So it's going to set a price that is going to clear both on the consumer side and the worker side because workers are getting a fixed reimbursement from that price. And the platform also chooses an assignment rule S . So that's the unique piece here. The assignment rule says what share of rides that show up will I assign to the kind of high preference workers versus the low preference workers? So those are the two major levers. And now, let's think about what this does, how this transforms into a wage rate. So here's kind of a rewriting terms of the model. So wages are going to arise from consumers arriving to the market. This is Q coming from the demand curve times the probability that any consumer is assigned to any driver.

So that's going to be a combination of the likelihood that a chosen consumer gets assigned to a high type driver or low type driver as well as the total number of drivers in the market. So S divided by N . So that's the likelihood of any driver meeting one consumer. And what do they get when they serve their ride? They're going to get the reimbursed share of the price. So this is P times one minus R . And this is going to be an 80% of the ride price.

Okay. Now, the model has a lot more to it. I am simplifying this. So when we offer kind of different wage schedules to drivers, they're going to accept or not accept, and the platform will see that over time and be able to learn about the underlying opportunity cost of drivers. And that becomes the basis for learning how to set these assignment rules optimally because we're going to set them such that we get the employment schedule that we want as the platform. Okay.

Here's some other potential issues to consider. The platform may want to keep the price signal to itself and not offer drivers a chance to see how much they value their trips. So the assignment rules are going to be opaque. And so we can think of reasons why that might be better, keep that information in the hands of the platform.

Second, they might introduce a stronger intertemporal commitment in the sense that if I assign you a series of trips during a busy time when you're likely to quit, if I keep you busy, then I'm also keeping the drivers from facing the decision point. Do I want to stop now or not? So you could think of the assignment rules as imposing this kind of intertemporal commitment by giving them less down periods. And finally, they may be able to use these assignment rules to control match quality, which is kind of a set of features missing here.

For example, a trivial example, maybe I want to match ~~high~~ drivers with ~~high~~ riders, a ~~five~~ star driver or the ~~five~~ star rider. And we could think about reasons why that might be beneficial. I know I'm

food delivery, service quality is very important because I don't know if this happens to you, people steal your food, and you get your soup spilled over everywhere. So they care more about service quality, and they do include that into this assignment score. Yes.

Audience:

So I thought that was interesting. I was wondering, you clearly only have data on ~~one side~~ platform, but it seems like these assignment mechanisms could make it hard for drivers to kind of multi home. And so they can serve as a little bit like an exclusive contract. And just again, you don't have two firms, so this is more of a simulation exercise. But do you have any idea of... If you can say anything about that

Yanyou Chen:

First, very good point. Not for this project, that's actually what motivates our project. We were first

And we have a set of theoretical results and that we have a model which I think is really modeling non-competitive work from the ground up and is quite illustrating in terms of how the economics of non-competitive work. And we have some sort of stylized results that show that non-competitive work can, in particular, if they become widespread, really erode competition in labor market and probably suppress wages. And then we add, then, we'll put some numbers on it.

So to the first part, it's very much a modeling exercise. And so it's taking off the shelf, again, the canonical model of competition in the labor market for workers via posted wages, which is the Burdett-Mortensen model. There's been many, many features with the Burdett-Mortensen model, which many of you might be a bit more familiar with. And then we'll introduce a couple of new features that, as I think, are useful if you think of modern competition applications in the labor market.

outside firms, to the outside employers, even those that might not have the capacity to write this type of contracts, to enforce these type of contracts and so forth.

We show one rationale why even based on purely utilitarian efficiency grounds, you might want to ban this stuff which induces misallocation of workers across firms. Then I'll talk a bit about welfare, and the welfare results are a bit ambiguous in the following sense. Worker turnover in these models is something that's actually quite inefficient because it generates a lot of churn. So competition generates churn, and that generates a lot of inefficient hiring cost, turnover cost. So an upside of ~~competition~~ is that they sort of economize on that. So, I'll get back to that.

Then we'll do something quantitative, a sort of interest a little bit in the comparative statics. Where would a ban really lift up wages? Where would it maybe not? So what we find is that there's some interplay between market concentration and ~~competition~~. So banning ~~competition~~ lifts wages, in particular, in concentrated labor markets. It increases wages when turnover costs are high because that's when frictions are high, and so that's when ~~ant~~ competitive conduct and rent extraction motives can really shift rent. Then we'll talk a bit about the role of the product demand.

As you can see, if the product market is highly elastic, then firms just can't pass through any of the rising cost from an increase in turnover to consumers and that basically has all the gains to workers evaporate. Then I have, just given that I come maybe from a bit of a different angle, a couple of things that might be sounding a bit pedagogical in the sense that some things we learned is that if you operate in this environment, you have to be a bit careful with some of the things we measure and we interpret if we live sort of in other settings. In particular, I'll talk a bunch about how to interpret quit elasticities, retention, elasticities, things people commonly measure now as a measure of labor market competition. Anyway, I'll get back to that.

Okay. So I don't want to talk much about the literature, but just to avoid any confusion, so I'm using dynamic monopsony in ~~ple commonly measure now as a measure of. to wre0 1 7(rke)3(e the 62 415.t 1 0 yrsre~~

how well-off our workers, because here, it'll move in opposite directions. If churn is up, markdowns are [inaudible 01:21:27]. The final other thing I want to say on this is there's now a whole industry of paper set. There are people out there that measure quit elasticities, retention elasticities, these type of objects, and then draw conclusions on how well workers are off, how competitive the labor market is, and so forth. Where does this logic come from? It sort of comes from a notion that, in a highly competitive labor market, firms basically have to pay the marginal product, and you can't really deviate from that. So the quit elasticity or retention elasticity is really high.

Now, think of what happens in these type of models if competition completely unravels and collapses. In the limit, everyone is just paying the outside option. Everyone is paying workers the flow value of unemployment. So workers are in a really bad place. The labor market looks terrible. What's the quit elasticity? It's infinite because everyone pays the same wage. The outside option is the flow value for unemployment. So in this type of models, these measures can be really misleading measures for the type of things people make of them. So that's another sort of maybe some pedagogical comment.

In general, there is a clean neoclassical mapping between labor supply elasticity, these quit elasticity,

compensating differential for the fact that you basically sign your life away and you commit not to move and transition to better opportunities. Okay?

As you can already see that just in terms of wage differentials, that the wage is not very informative about values or how attractive the job is or how much damage ~~acompete~~ does. Okay? So I can't really... This is going to be hard without me pointing to the right spot, but you have on the left the wage offer distribution, and the blue, it's the equilibrium wage offer distribution without ~~acompetes~~. It's just a uniform distribution. Okay? So then I'm going to give the first firm in the labor market access to non-competes. That's the red line. So what you can see here is that that firm posts a mass of jobs with pretty attractive pay. That pay is kind of in the middle of the distribution. Okay?

So now, you might say, "Well, that looks pretty good." But then on the right, you have values. So not just wages but the full forward-looking part, which includes the option value to climb towards more attractive pay. So there you can see or almost see that these jobs piling on, there's mass now at the very bottom. Okay? That's that piece of mass on the red line on the right. These jobs that look pretty decent in wage space look pretty terrible in terms of value space because they're, again, the least attractive jobs in town. Then you can basically show that from that, you get pretty strong spillovers to the rest of the market because these guys are basically no longer competing. They're sitting at the bottom of the ladder.

As a consequence, there's sort of less competition on the interior and all the other firms start reducing pay, and then reservation wages start falling, and so forth. So you get this general equilibrium effect and things start deteriorating for workers. What's the limit? The limit, you can sort of see it. I don't have to pay you a compensating differential anymore. Everyone signs ~~acompete~~. Everyone gets exactly B. Nobody has any incentive to deviate and offer something more attractive because there's nobody to be poached because everybody's under ~~acompete~~. So think completely unravel and you restore the Diamond equilibrium that you've sort of undone with the job ladder competition forces prior.

Okay. So I guess that's what this slide is saying. I mean it's obviously stylized, but it's just to point out that once these technologies become or contracts become widespread, really has the potential to unravel competition in these type of markets. We have a bit more on welfare. So there's a couple of things there. The first is that what you can see is that, these firms in the basic model, all have the same decreasing returns to scale technology. So you don't want them to have different size because that means misallocation because that means there must be a differential margin product of labor. So when

but basically, what we're doing is we're doing two model validation calibration type exercises with the

that there's heterogeneity in the prevalence of noncompetes. That's an aspect I want to kind of come back to as well. So what does this paper capture? Wage effects? It's very well set up to capture that these noncompetes are dampening down wages. In terms of misallocation, yes and no, I would say. I mean the baseline model, everything is homogeneous. All these, we call it a job ladder. Really, we mean a wage ladder. They're all doing the same thing in the same kind of firms.

We can throw some heterogeneity on the firms in terms of productivity, and that allows for some notions of misallocation. But I'd rather be at a job that's a five minute walk from my house than an hour away from my house. We don't capture any of that or any of the things that normally live inside the size and other Greek letters you guys like to use so much. So I mean, I'm being glib, but I don't quite know how to think about that, how seriously to take that misallocation. What about the good of the noncompetes? Gregor, in the paper says, "Well, a lot of this is happening at low jobs, the kind of famously Wendy's non

Gregor Jarosch:

Yeah, I should first say thanks a lot for discussion. That was really great. So I don't know off the top of my head, but in the baseline, much of the pay raise gets passed through, so it's not much less than 4%. That's right.

Speaker 6:

Speaker 8:

Yeah. One of the arguments made in favor of ~~no~~ competes, at least in very specific situations, specific roles and so on is that, in the long term, if ~~no~~ competes weren't allowed, it would disincentivize

think of the work by Ali Akoglu. Cases trying to understand bargaining between health insurance providers and hospitals has really tried to think about the tools that you need to analyze concentrated markets in both on the upstream and downstream. And the goal of this project is to think, can those tools be used to analyze the employment market for teachers?

And I'd also add, unions are prevalent in a lot of places in the world, so like 40% of Quebec's labor force is unionized, or all of Germany uses sectoral bargaining. And there're not many unions in the United States, but where you do find unions are actually in the public sector. So this is the place where we see them the most. So with that in hand, what does this paper do? We have got very detailed granular data on teachers in schools in Pennsylvania. If you're following the last discussion, and Heski was asking about the things that IO people do with Greek letters. This is like a representative of that approach. We're going to use this Nash-Nash bargaining model that's been used extensively in IO to understand markets where there's power on both sides and apply it to the setting of collective bargaining for unions.

And I should just stress, a lot of that work was set up to understand different type of mergers that were going to be analyzed, that people in this room have worked on. We're going to apply that to the labor market. So a lot of the goal of this paper is to try to port over some of that work to the labor problem. And we're going to use this model to think about what's the efficiency of unions, what are the outcomes if you don't have unions, what are the socially efficient outcomes? So try to do kind of a welfare analysis in this context.

All right, so let me just give you some setting. So do the IO thing of diving into the details. So we're going to look in Pennsylvania, there's about 500 regular school districts. There's also a number of charter schools, but most people work in a regular school district. And the way that regular school districts set wages is through a collective bargaining process with teachers' unions. And more importantly, each school district has a local teacher union. So the bargaining is going to happen at the school district level on both ends. The second thing is, to the extent that monopsony power is distortionary, it's really important that when you hire an additional worker and you have to raise the wage to hire that worker, you also have to pay everybody else more. And to make that very clean, you need something like uniform wage schedules.

If every worker gets paid a different amount of money, it's a little bit harder to understand how that distortion is going to work. And here, the setting is going to be simple. The amount you get paid as a teacher depends on, do you have a master's degree, and do you have up to 12 years of experience in teaching, more or less? So a very simple kind of salary schedule context. And the other thing to say about Pennsylvania, say for instance North Carolina, where wages are pretty much even across school districts, is there's a lot of variation. So Lower Merion outside Philadelphia has average wages of about 100,000 in 2016, versus other school districts like North Star have wages under 50,000. And even school districts that are close to each other like Philly versus Lower Merion that kind of border each other, have

And so the efficient point would be at a point like A. Oh, sorry, I'll go back. I guess this wouldn't work. The efficient point would be at a point like A, but the monopolist is going to basically choose a lower wage, which is going to be at the intersection of this blue line, the marginal benefit of labor and this kind of dotted marginal factor cost. So the monopolist is going to choose fewer workers and a lower wage. Now let's try to think of what happens if there's say a collective bargaining agreement. In this context, it can choose a different wage. So it doesn't have to choose the monopolist wage. It can choose a wage that's higher or lower than the market clearing wage. So here I've just drawn it where they choose a wage that's above the market clearing wage. Here, there's going to be a different issue, which is, is the number of worker hired going to be on the labor supply curve, or is it going to be on the labor demand curve, right?

Because there's a difference between labor supply and demand. I'm going to use something that in the literature they call the Medoff union model, which is, we set wages, but then employers get to hire the number of workers that they want to hire. So in this context, they're going to hire the number of

And if you don't have this mechanism while they just didn't get an offer in that good school district, then

background. All right. Okay. And let me move on to the relevant counterfactuals we're thinking about. So the first one is, what does the world with posted wages look like?

What does the world with the social planner or as people at Chicago told me the intersection of labor supply and labor demand, what does those look like? And then what do wages look like under the Nash bargaining model that we've estimated? Okay. And so I'm going to be presenting kind of simulations for the entire state of Pennsylvania. I'll be showing you weighted median wages just indicating that outliers are still an issue in this 0.000ion of labor

Allan CollardWexler:

Yeah, I think it's a great question. I mean even broader, the proposals having statewide bargaining versus not are kind of interesting. There's some issues with how you set up the objective function of school district when you combine them that we haven't figured out yet is slowing us down now.

Allan CollardWexler:

So we went through the collective bargaining agreements just to see the timing. So they're not negotiated at the same time. As you might expect they're represented by the same umbrella organizations. There's two points. So one is should care about other unions. So what is the objective payoff in the Nash bargaining of unions and can we play with that? I think that's where it happens. On your previous point, the Horn Wolinsky paper that gets used for Nash bargaining actually also has a union paper published the same year. So that idea's been around for a while.

Speaker 4:

All right, one more question.

Speaker 12:

Hey, I was just curious about teacher quality, and since you have this very, very low number of offers that comes out of the estimation I think is there some sorting going on in school district on that front?

Allan CollardWexler:

Yeah. We have this movement not being correlated with rank. There's some other evidence in this literature about value added not being correlated with pay, but the idea that all employers are seeing something that I don't see that's not reflected in value added, I'm not sure what are the ways to either validate or disprove that, but that's the big outside explanation here.

Speaker 4:

All right. Thank you so much, Alan. We now will break for lunch which will be in the back corner there and we will resume at 1:00 P.M.

Sam Kleiner:

Okay, welcome back everyone. I hope everyone's ~~feel~~ and ready to go for the next session. My name's Sam Kleiner I'm a staff economist at the FTC and one of the ~~organizers~~ along with Viola. I wanted to take a few moments to reiterate our thanks to all of the folks who made this conference possible. It takes a lot of people doing work behind the scenes. So I wanted to first thank all the folks in BEU who helped us out with the selection and just the organizing duties. Also wanted to thank our two points of contact from the FTC event planners, Pinar Gezgec and Bruce Jennings who were the main points of contact for the production team. And I specifically want to call out Stephanie Aaron. Is Stephanie in the room somewhere?

If she's not, give her a ~~shoot~~ if you see her. She's the one writing all the emails to you. I'm sure you've seen her name on the emails. She is I think the true of this conference who helped put everything together. She has amazing attention to detail with all aspects of the conference down to every last detail from the call to papers to checking all of the slides. There she is right there on the door. So she put everything together. So thank you, Stephanie. So with all that said, let's get back to some more research. This first paper is going to be presented by Andrey Simonov, the Gary Winnick and Martin Granoff Associate Professor of Business at Columbia Business School who's going to be presenting the paper, What Makes Players Pay: An Empirical Investigation ~~Game~~ Lotteries. It's going to be discussed by Fahlha Ibnat, a staff economist at the FTC. So with that I give you Andrey.

This is not exactly how the game looks. We asked an artist to give it to us because we couldn't show you the actual game, but it's very typical. In the middle you have this progression stage to stage, and on the right have a loot box, this lottery which a person can decide to open. Notice that by Japanese regulation they're required to make probabilities of these different quality very visible. Now once a get to model, we'll assume people know these probabilities. We can have a discussion about what else we could do here. We'll assume people have rational expectations and no probabilities. Okay, so that's our context. We have complete access to data. As I mentioned, it's a two and a half million people.

We see actions of play, opening loot boxes, outcomes, inventories, the currency stock, a lot of different things. Basically the same thing as the company says. A couple of summary stats I'll highlight. It's a big table. So again I wish I could point, but the blue square on top gives you on average how many main stage, which is what we'll focus on people play, and average is 38. An average opens around eight rare loot boxes. This will be this loot boxes you pay for. You actually need to spend money either in

definition is random. So sometimes you get good news, sometimes you get bad news. We can see if once you get good news or higher quality item, you are less likely to go and use the loot box more and now you'll go and play the game. Why? If it's functional value you should be more likely to go and to use this item and enjoy winning the game and progressing in the game more. Okay. So we'll do this model free evidence with a simple IV regression. We'll look at the probability to open loot box after opening loot box right now. So open it again.

We'll instrument the quality of the inventory, this R , with the outcome of the loot box, which you just got, which we know is random. In this paper we'll summarize the quality of the inventory by the quality by rarity of top four items you have because you can choose up to four items when you play the stage. And then we include user fixed effects as well as stage by your inventory time T minus one is fixed effects. If there was no missing data measurement error we didn't need IV, we control inventory T minus one and just can do a less. Sometimes we have missing data so we'll also use IV for rarity. Okay. So we run this, what do we get? For regular players, ~~males~~, if they have a one x extra rarity item in the

counterfactual on product design where we'll increase or decrease the difficulty of the game and see which players, how much revenue we get and how much engagement we get from players.

It's interesting because the current design of the game, it turns out nicely balances revenue. The company gets from whales, where they get most of the money from engagement, where they get from regular players. And so if you try to put a value on customers who wouldn't pay to the company right away, that's the one way you can do it. Why would you want to have paying customers? Because popularity of the game will allow you to promote it and to make it much easier to sell or to distribute among these paying players.

And so the final thing I wanted to show is now trying to decompose as different parts of utility into loot box and gameplay utility on the different constraints of what the firm does. So here we have welfare on the slide. Just for the record, welfare here means any utility we get in the model. This can be preferences which we think are respectable from welfare perspective or things like addiction. So we can try to correlationally remove some parts, but we don't have a good way to pin down one or another. So taking all of that preferences as surplus.

In the baseline under the current design, the firm gets around 7% of surplus. That's our revenues. Players overall get 6% of surplus from loot boxes and 68% from playing the game itself. For whales, which are in column C here, this is different. So they get around one third from loot box in terms of their surplus and two thirds from playing the game. So that's the first bars in these three pictures.

The second bar is if we put a full blanket ban on loot box without changing anything else about design. Then the consumer surplus, producer surplus goes down by definition, but consumer surplus goes down by 25% because of the complementarity. For these regular players, that's where this functional value comes in. That's why it will not be great.

We can put a ban on paid loot boxes and then regular players recover almost all of their surplus. And finally we can think about different counterfactuals where we put spending limits, which allows us to recover all surplus for regular players and more surplus for whales. So even if we put the spending caps at \$100 per person, even this high spending players already get 84% of the surplus they got before. And

Viola Chen:

These are, yeah.

Speaker 13:

There's a lot. Do you guys have any questions about [inaudible 03:04:22]?

Viola Chen:

Great. Hi everyone, my name is Fabliha. I'm a staff economist here at the FTC Bureau of Economics and I want to thank the authors for writing and presenting this paper. It was very interesting and also incredibly relevant for some of the work we're doing at the FTC in trying to regulate loot boxes.

So I'll start with that, why this is a very useful paper from a consumer protection perspective. Regulating loot boxes is a relatively new consumer protection issues at the FTC. And a key challenge that we face when we're determining if a loot box system is unfair is determining both injury from loot boxes to players as well as any countervailing benefits to players from loot boxes. These are difficult things to estimate and they require understanding, in part, how players view and interact with loot boxes.

So this paper provides a very useful framework to start thinking about this question, as well as convincing empirical evidence for how whales and ~~whales~~ whales tastes for loot boxes might be split

Other than that, I just had two other more minor questions about other components of the analysis. I'll skip the first one because the authors didn't really go into it in this particular presentation and it might be a bit too much detail. But I'll go back to the spending cap counterfactual. The authors note that this counterfactual is simulated in a stylized way where players are actually myopic regarding their budget restriction. So they don't anticipate that they'll hit a spending limit.

I'm wondering why this myopia didn't really affect players' ability to fully recover their play utility by reducing their utility from winning. So in that previous graph, that blue bar, I don't know if I can... Yeah, the blue bar, I would expect to be a bit lower for maybe some of the lower spending caps, but I see that it's about the same compared to baseline. I'm wondering why that is because if players are myopic, I would expect them to not be able to strategically time their loot box openings or take advantage of the nonlinear pricing of in-game currency in order to buy more loot boxes and open more loot boxes within that budget restriction.

I'll finish by just highlighting three other future research ideas that I think would be particularly useful for policy. I know that these are all beyond the scope of this particular paper, but these are some of the big questions we've been grappling with at the FTC as we try to think about how to regulate loot boxes.

So for one, it would be useful to understand the effect of confusing or unclear loot box odds disclosures on player behavior. Video game companies tend to argue that odds disclosures that may seem convoluted to the layperson are actually very understandable to a frequent player or a whale. But there's also anecdotal evidence that players have been confused and overspend because of that confusion. So this is a useful empirical question that would help guide policy.

The second question that would be interesting is investigating the effect of disclosing the average cost of obtaining a desired loot box prize on player behavior. As regulators, we're worried that players don't actually understand how much money they're going to ultimately need to spend in order to get a five star character or the desired character. And regulators say that these average costs are actually very difficult to calculate because it's hard to define the average player. And there's complexity when you start thinking about the nonlinear pricing of in-game currencies. So again, another empirical question that would be useful.

And finally, it would be useful to understand differences in player behavior by age group. As the authors mentioned, a lot of these players are children and we believe that they're particularly susceptible to some of the harmful behaviors that can result from using loot boxes. So it would be useful to understand how they interact with loot boxes. But that's it and I want to thank the authors again for writing this paper. It's very interesting and very relevant for our work here. Thank you.

Viola:

We have time for a few questions.

Speaker 14:

It is really an interesting paper. I wonder in the counterfactual spending caps, that means the game developer will receive less revenue, right. If you're thinking this may reduce their ability to elevate on the game later on, that might hurt other players to enjoy the game. So I wonder how you think about that trade-off.

Andrey Simonov:

Yeah, no, that's a great question. So first of all, let me thank for the discussion. This was excellent. And we can also connect offline because I have other follow up questions about this. This is great. So on this

question about... Yeah, so in the counterfactual we didn't allow them to make any adjustments so it's partial equilibrium because they'll adjust the game in some way, especially if we don't allow any paid loot boxes, then they probably wouldn't produce the game to begin with. So it's a good question how they'll adjust.

One thing I'll say about this product and this game, if you look at the activity over the four years that we have the data, there was a big spike in the first maybe 3, 6, 12 months and then sometimes they initially update, but by the end... It took another two or three or four years for them to discontinue it even though there were few players.

So I would think that even this, they are doing this game in one or one and a half year to recover as much costs as they invested to get some revenues. And then the marginal costs of just keeping the game are not very high and so they'll probably discontinue it either way. So I don't think it'll be big adjustment. Now if we put really a strong spending cap, that will be a different story because then it'll change how they do it. Thanks.

Speaker 15:

So full disclosure, I'm the parent of a ~~ten~~ year-old and twelve-year-old boy so you had me at FIFA. I want to come back to the last bullet point. I mean I know you don't have, I suspect you don't have or you would've used them, demographics of the users explicitly, but there's things that you can easily imagine as correlates. Like school holidays versus ~~school~~ school holidays might have a much bigger impact on school children than adults. I wonder if you played with that at all.

Andrey Simonov:

Yeah, that's great. So we did, not as successful as we were hoping we will. The only demographic in the data we observe is a breakdown by the operation system payment type, so iPhone, Android so it's not great. Using my course's knowledge, we try to think about school holidays and timing of the day where school ends, the school starts. I think in Japan there was regulation that if it's a minor who uses the phone, they cannot spend more than a hundred dollars per month. I'm not sure, but so we tried to look for these patterns. We didn't find as much. If there were ideas of how we can try to recover this would be great because we do have the timing data. We just don't know who those guys are. But yeah, thanks. That's great.

Speaker 16:

Yeah, so a really interesting paper. We went pretty fast through the structural model, so I am not sure I caught everything. But I was wondering, so in your model players, the whales ~~whaddies~~ know they're addicts and you could envision a quite likely scenario there's sort of multiple cells or internalities where addicts, they're not myopic, they're forward ~~looking~~ looking but they don't internalize the addiction part, and whether this would matter for the results.

Andrey Simonov:

No, that's a great question. So yes, so we assume they fully... In some ways they assume they know all their preference. Everything is part of the utility, everything is part of the choice and they are forward looking. I didn't talk about transition probabilities, but we think that they correctly anticipate the odds, the transition. So all of that, it's a more standard single agent dynamic model and estimation.

We tried to go more into this decomposing the welfare respectable preferences from other. We initially thought of state dependence as some proxy for this, but thinking more about the context, this repeated

actions of lotteries, there can be a lot of functional reasons why if you didn't get what you want, you'll go and do it right away, right immediately.

So one thing actually we observe in the data which I can mention is that the repeated action on open another loot box, it takes on average like five or seven seconds to open one more of the lotteries. The action between the action of loot box and play or play versus play, even once people finish the play, then it's more like one minute. So there is this repeated openings which will be consistent with some addiction or impulsive consumption. They happen very quickly.

But then again, it's hard to interpret because maybe these things happen for functional reasons, not because they're addicted. Maybe, I know I need something else and I'll quickly click there. So in the paper we try to stay away from interpreting one way or another, but ideally, if we had some identification to separate it out would be great. And we're happy to talk more about this offline.

Viola:

All right, thank you so much, Andrey.

Andrey Simonov:

Thank you.

Viola:

And our next presenter is Mark Shepard. He will be presenting on adverse selection and unnatural monopoly in insurance markets.

Mark Shepard:

Thanks so much, Viola and thanks for having me at the conference. It's really a pleasure to be here. Thanks to the organizers for including us on the program. Let's see if I know this. This is joint work with Ed Kong who's an MPhD student at Harvard as well as my colleague, Tim Layden, who's now moved to University of Virginia.

markets in a variety of health insurance settings are highly concentrated by standard metrics. And this is particularly severe in the Obamacare or Affordable Care Act marketplaces that were set up to cover the uninsured starting in 2014.

Here's data on the Obamacare markets from a couple of years ago showing the number of insurers participating in each market. And you can see across county markets, it shows the number of insurers, number of counties where there's one, two or three or more firms. And when the scale is one, two, three or more, you know you're going to be in trouble with very limited competition. In fact, at that time about half of counties comprising 20% of the population had just one or two competitors, monopoly or duopoly, and 24 whole states had three or fewer competitors.

What's the four firm ratio when there's only three firms? I'm not sure. I guess it's just the top three. It's the three firm concentration ratio. This is a problem. That's our basic point is that low competition is severe in many insurance contexts, but particularly in the Obamacare exchanges. And what explains this? It is a motivating fact. We wanted to understand why robust insurance competition is so difficult to sustain in many settings, particularly in the Affordable Care Act markets. And of course there are going to be some standard factors that are relevant here as well as in other settings, regulatory barriers to entry, fixed or sunk costs of entry, political factors that were particularly important in Obamacare. These are all important.

We want to though argue that there may be other complementary features, an additional feature that's been missed in prior literature and that we think is important and concerning is a fundamental feature of insurance markets, and that's the classic market failure of adverse selection. So we're going to argue there's a connection between limited competition and adverse selection.

What is adverse selection? Classic insurance market failure. It's one of the key things that makes insurance markets different. It's the property that sicker people, those who are higher risk tend to have higher demand for various types of insurance. And there's asymmetric information or unpriced cost heterogeneity that leads to a variety of problems.

Now typically, adverse selection has been associated in classic work with a couple of things. First, it's with markets not functioning or market unraveling or sometimes it's called unraveling of trade. Think about the classic Akerlof model. We can't get trade in used cars because there's asymmetric information and only low quality cars are sold. Okay, so that's one market failure.

Another classic market failure with adverse selection more associated with the Rothschild and Stiglitz classic model is about quality unraveling. So all firms degrade their quality to try to avoid high Consumers who again are unprofitable because we can't price discriminate to cover their higher costs.

We want to argue and ask whether there may be another market failure associated with adverse selection, which is that in some settings where quality is heavily regulated and trade is insured via mandates and subsidies, could adverse selection also be a barrier to robust firm entry and competition in these types of markets? And that's the basic argument of this paper.

We're going to argue for a mechanism by which adverse selection, classic market failure may lead to robust competition. It'll be a new implication relative to prior literature for what adverse selection can do as a barrier to robust entry. So how does this work? The goal of this paper is to teach you a concept. So let me summarize it on this slide and then I'll show you how we build it up in the model and the empirical work that we do going forward.

Well, here's the key insight. Think about a market, I'll give you a specific example in a little bit, in which firms are differentiated and they're competing on prices. So they've entered the market and they're competing in standard differentiated Bertrand style on prices. We're going to argue that adverse selection can create incentives for something that looks like a race to the bottom in prices.

good. The ideas go beyond just the simple model. Set up, we're thinking about insurance market where there's potential firms J , who engage in a stylized ~~stage~~ entry game. Obviously, the real world is dynamic. There's many other complexities, but this is a simple model that's been a workhorse in the IO literature to convey ideas, and that's what we're trying to do here.

In stage one, entry; firms simultaneously decide are we going to enter the game? Stage two, the entrants engage in standard Nash price competition. Each insurer will have a single fixed contract. They're differentiated. We're going to treat that as exogenous for the model, although I think it's something in extensions. We may want to think more about endogenous differentiation. But importantly, we want to think about general horizontal differentiation among firms.

So often in adverse selection papers, it's always been about quality. It's always been about quality H plan and a low quality L plan. And the high quality plan gets screwed over by getting so many ~~sick~~ consumers. That's not what we're going to be about. We're going to be about plans that may be horizontally differentiated, and yet we're going to argue on things like, by the way, hospital networks, which will be the key example. So if you think about firms that differ on which specific providers they cover, that's a horizontal differentiation. Yet we're going to argue that adverse selection is still important.

Consumers will vary both in preferences θ 's

for different reasons. On the left, it's because if I as a firm cut my price, I bring in relatively low marginal consumers. I'm moving down a downward sloping cost curve because of risk selection. On the right, it's because I have big fixed costs. And as I get more consumers, I spread those fixed costs over more people.

But in both cases you get this downward slope, you get this wedge, and those two things, maybe not surprisingly, will mean that both will behave similarly. Adverse selection, just like fixed costs in our classic theory, will serve as a reason why a market can support fewer firms than otherwise. And if it's very severe, if adverse selection is very severe, just like if it's fixed costs or extremely high, you may have

margins for firms in equilibrium? Their price minus average total cost will equal the Lerner markup, minus the wedge between average and marginal costs due to adverse selection, minus fixed costs per consumer. And this is the point that we're making in the paper, that both adverse selection and fixed cost per consumer enter in a very similar way to the total net profit margins, which determines whether firms are willing to participate. So in both cases you're sort of getting lower markups, which therefore may lead fewer firms to be willing to participate.

We will, just to be more specific, we'll use a difference-in-difference design that compares how demand changes for high and low income consumers, or relatively higher and lower income consumers in this market, where you can think about here's the five plans in the market over time. They vary their prices over time just due to normal price competition. And higher income consumers face those price changes. And we can see, okay, if a plan raises its price, does it get sicker consumers? Does its average cost go up? We can compare that against a control group of the below poverty enrollees who have the same plan menu but don't pay prices and that can net out any unobserved demand or cost shocks.

Here's, just briefly, the descriptive evidence. Here's our first stage. We're going to think about plans that raise their price at time zero versus those that lower their price at time zero. And it's all compared to that below poverty group as a control group that would net out any demand shocks. In practice, there's very little trends for the below poverty control group, so you can think of this as all coming, all the variation coming from the treatment group. Here's the premium changes. They're relatively small swings

multinomial logit. We're using our observed micro data to see how people choose among health plans. We'll have heterogeneity in price sensitivity coefficients by risk to capture that adverse selection. Cost comes from individual risk as well as plan-specific cost effects, so we can capture if certain plans are lower costs than others. We see a little bit of that, although mostly there's not a huge amount of variation there.

And then equilibrium, we're just going to put it into our two-stage entry game. Stage two is pricing conditional entry and stage one will be entrance among the set of firms who we saw in the Massachusetts market. Let me actually skip this in the interest of time. We fit the model pretty well in terms of when we simulate, use our model demand and cost estimates to simulate out those difference-in-differences plots that we showed you in the reduced form. We can fit those quite well. That's natural because we were basically fitting the model off of that type of variation, but it's comforting to see that you fit it well in sample.

What are we going to do in the counterfactuals? We're going to think about what happens if you take away some of these policies and we'll fit our two-stage entry game, I said that before. Just a couple additional details of what we do with the counterfactuals. Our potential entrance is always a concern in

simulation. It's quite a bit stronger than the actual risk adjustment that was used in the Massachusetts data, which was closer to about 0.3 or 0.4.

Finally, price floors, again the controversial policy, but it works pretty well and it's pretty simple to

healthier patients are also more price sensitive... Actually, less price sensitive, then what you will see is that the firms, the insurers have incentive to price lower more aggressively, in order to attract those healthier patients because they are lower cost to treat and they're more price sensitive.

And now, this is a point that actually has been acknowledged and has been pointed out in some previous papers, that adverse selections can lead to more aggressive price competitions. But the authors push this further and argue that this fact, these implications could actually lead to a lower profit. And as a result in equilibrium, you might end up having fewer firms to be in the market because the profits are not enough to cover the fixed cost. And in some extreme cases, if the adverse selections issue is serious enough, you may see that there's only one insurance provider that can survive. And when that happens, that insurance survivor, that only survivor, is going to charge a monopoly price and that will lead to a very high price. And so the authors propose that to use a price floor, a very ~~controversial~~ idea, to correct this problem. And because the price floor is going to limit competition and that might actually lead to more firms in this market.

Now what I want to think about is, this is all good. I really enjoy, appreciate these ideas. But I want to just deviate from the theoretical models a bit and think about what if the insurance providers can vertically differentiate, in the sense that maybe one insurance provider will provide a better quality of

Now, I'm not saying that in this paper the authors need to really explicitly do it, because obviously this is very challenging problem and it's a high dimensional optimization problem. It's very difficult to do. But perhaps one can impose some structures to the problem to simplify it. For example, maybe allowing for a few preset network configurations in order to address this. And I think I can share more thoughts with Mark.

And so now other detailed comments that I have about the monopoly situations that I think Mark pointed out earlier, one thing I'm wondering is to what degree that we actually see the market where we have monopoly, how high the prices really are? In the paper, I think at this point I haven't seen any figures that show that. And perhaps the authors can provide some evidence on this. And the one thing that they argue is that theoretically they don't expect, when it's a monopoly situation, they think it's going to stay that way and the price is going to be high. But if it turns out that the price is not as high as what the theoretical model predicts, I wonder could be happening, right?

One possibility is that maybe the threat of entry is disciplining the monopoly pricing. And the authors actually acknowledge this, but they argue that it's not going to be profitable for potential entrants to enter the market. But what I'm thinking is that, it is not entirely clear to me because theoretically, even though if a potential entrant sees that this is very profitable, and if they enter the market, they may have something like 500 odds that they can kick out the incumbent. And now, of course, the incumbent may actually have some absolute advantages. But still, if the profit is very high, it might still induce the potential entrant to try to enter the market, and in the sense that if the expected profit makes sense.

What could explain that? I don't know. But potentially it could be that, the potential entrants. We may not really know exactly what the demand is, looks like. And so in some sense if you don't price it too high, you might have preserve some of the [inaudible 03:59:46] information, and that could be one way to motivate why the monopoly market may not be pricing too high.

But anyway, these are some conjectures, and I have some other comments. I'd like to see more evidence that the adverse selection story could be really happening. And I do believe that this, to a large degree, it's going on in this market. But it would be nice to see that, for example, when you look across the market, is it the case that we see the more insurers in the market that we have more insurers, we have less patient heterogeneity in terms of their health status? That's a measure of how serious the adverse selection problem is. And on that hand, I also want to see some evidence about maybe potential sorting in equilibrium. And for example, when you look across the market, to what degree we see the insurance network differentiate among themselves in terms of high quality or low quality and so on and so forth.

And now I have natural questions that when I read the paper, it started off by using ACA to motivate this study. But at the end, when it comes to the empirical exercise, they used only the ComCare data, Massachusetts data, to calibrate the model, estimate the model. But Mark also pointed out that ACA data actually is not as detailed as the ComCare data. And so what I think the authors could do, is maybe use, discuss a bit more about to what degree other markets are similar to the ComCare market, the Massachusetts market? And to what degree they can argue [inaudible 04:01:34] that difference, how they can deviate the parameters a bit and say that the result is still robust, or how much confidence policymakers should have about those results.

And so to conclude, I think this is a great paper. I think it's a very interesting paper that generates new insight about the implications of adverse selections on firms' entry. And I learned a lot and I'm over time, so I'm just going to stop here. I encourage you to read this paper. It's really ~~with~~. Thank you.

Mark Shepard:

Thank you so much Andrew, for those really helpful comments. And we've talked a little bit offline, but you're absolutely right that I think we are conveying insight here, but we need to do more work to say how relevant is this insight for what's happening in the ACA? And I took away that very clearly. Thank you so much. Other questions?

Speaker 16:

I had a question. So you are focusing on adverse selection on the intensive margin?

Mark Shepard:

Yes.

Speaker 16:

So if a plan lowers its premium, it's going to get healthy people, but they're all being stolen from the plan's rivals. I'm wondering if your conclusions would hold, do you think, if the concern is adverse selection on the extensive margin? There's uninsured people that the plan could get into the market by offering a lower premium. And in particular, I'm concerned if you have a price floor, that you may be preventing that type of selection if that's an important issue.

Mark Shepard:

Yeah, it's a great question. In practice in the ACA, you have what are called price link subsidies, which ensure that the extensive margin is in some sense shut down. For subsidized consumers, which are 90% of consumers, the post-subsidy price for at least one of the linked plans is fixed by law regardless of what intensive margin price competition happens. I think in practice because of that, it's a little bit of an interesting institution, but very important, most of the competition is on the intensive margin. That would be one response.

In general, though, when you do have extensive margin, what you're going to get is... We've done a little bit of simulation work with this. It can help a bit if you have some extensive margin, if you're competing against an outside option where they're not strategically pricing against you, and you can lower your prices to bring in healthier people to the market overall, the selection problems are less strong. But it's an empirical question, is that decrease in the amount of race to the bottom in price competition big enough? It'll depend on how that occurs empirically. We think it's less relevant to the ACA, but it could be relevant in other settings where there's a fixed subsidy of some kind.

Speaker 16:

Very cool. Because none of the costs are sunk in this model, it seems like this is actually a rare opportunity to talk about something we don't really talk about very much, which is predatory pricing and whether there's, if I'm competing against you as an insurer and we do the thing where we cut prices to cream skim or cherry-pick, and that gets us so far, and then maybe I want to go just a little bit further because I know that's going to actually put you out of business. This is actually one of the rare occasions where I could actually conceive of that actually happening.

Speaker 16:

... other occasions where I could actually conceive of that actually happening. I wonder if that's something that you capture and something that you maybe see evidence for in the data.

Mark Shepard:

It's a great question. In some sense, that's what's going on here. Because here, directly by competing, by

One of those dimensions is a very immediate, very ~~short~~ effect displacement. You have a job and now you don't have it. There are other effects that are including wages that you might see, and there

like a Bellman equation where, again, you're thinking about how much a particular occupation should be

look at a set of outcomes. In the paper, we are looking at a much wider set of outcomes, but for this particular talk, I'm going to focus on essentially an aggregate measure of the career values in a particular location.

We can generate many different versions of this. But what I would like to understand is if we can focus on a commuting zone, a region at a particular time, and if we could think of all the potential occupations that are prevalent in this particular place, what does the weighted career values look like in this place. And if we can, look at two different time points. Do we see a change, do we see an improvement, or do we see a decline in the career values of these regions? And then we are going to try to understand, we are going to try to decompose how much of this is coming from robotization.

On the lefthand side, we can do an approximation through the labor market shares. That's something easy to, I think, go to. But on the righthand side, we need essentially an exposure to robot measure that we can plug in. Now here, in order to get to the robot exposure, we are not going to try to reinvent the wheel. There are already plenty of studies that are trying to get an approximation of the exposure to robots. But in a nutshell, what we are going to try to do here is to follow a paper, a series of papers by

robot per thousand workers, the average decline on the career value is about \$3,300, \$3,400. Now you might sa

Put differently, the negative effects of robotization on the local career market values are concentrated in these regions that are heavily focused on, or the labor market is heavily focused on manufacturing.

education, five to six to 10 years of education and more, which we do categorize even more experienced individuals.

For the more experienced individuals beyond 10 years, the numbers look very similar to 6 to 10. I just want to show you this very junior, as well as the next level of seniority individuals and their comparison. Here again, it seems like the numbers are similar, but especially when you look at the occupational career path effects differences, focusing on column four, you see that individuals who have a medium level of experience, 6 to 10, they seem to be more negatively impacted. This might be possibly because the type of occupations that they're going to move to are potentially more limited, and therefore they see harder times transitioning to those occupations.

One last thing I want to show before concluding. Another question that we might have, again, looking at individual transitions. Well, maybe they are not transitioning to other occupations, other jobs, but

transitions, we also see potentially some negative impact of robotization here in these key outcomes that we might want to focus on as well. With that, I'm going to conclude my talk. If you have any comments and questions that are going beyond this presentation, here is my email. Would love to hear from you. Thank you.

justified? I will show you experimental evidence that suggests that actually personalization helps consumers and smaller sellers on the platform. Then the next question becomes how will privacy restrictions that limit the ability of the platform to personalize impact different types of consumers, sellers, and the platform? I will show you evidence that suggests that privacy restrictions will primarily hurt more price-responsive consumers and smaller sellers on the platform. Even if we start accounting for consumers' privacy evaluations, still consumers will be negatively affected by the privacy regulation. Finally, in the last part of the paper, I start thinking what can platform do to mitigate the losses from privacy regulation.

In terms of methodology, to answer the first question about the impact of personalization, I use a two-year long field experiment that I ran together with Wayfair, where we randomly turned off personalization for a sample of consumers on the website. So this field experiment allows me to make causal statements about the impact of personalization. Then in the second part, to understand the impact of privacy regulations, because a lot of these privacy rules are just coming up, they're all in the future, so what I will do is the following. I will retrain the platform's actual personalization algorithm using lower quality data to mimic different types of regulations that are in the discussion right now. Afterwards, I will simulate how consumers respond to the counterfactual recommendations that would have been generated under lower quality data environment using the structural model of search. Finally, in the last part of the paper I propose a simple probabilistic identity recognition algorithm that can help platforms recognize consumers in a probabilistic way even when the platform is not sure who the consumer is exactly. Again, I will evaluate the probabilistic algorithm using the structural model of search.

In terms of the data, I have full access to Wayfair's data where what I observe is the device and the browsers that consumers use when they arrive on the platform. I know the source of traffic, meaning whether their consumers arrive directly to Wayfair.com or did they arrive from advertising channel. I know the product rankings that are shown at each consumer page load. Wayfair collects very detailed pixel-level click stream data, which means that all history of consumer actions such as clicks, scrolling behavior, zooming in on an image, zooming out, et cetera, everything is tracked. Finally, I observed the final purchase decisions and product returns, if any. So basically on the consumer side, I know what consumers see on the website, what they do on the website, and finally what they purchase or return. Finally, in the last part, in terms of the supply side, what I observe is the daily retail price and seller wholesale prices which allows me to calculate seller revenue and platform revenue and profit outcomes.

I hope the empirical setting is clear, and now I want to switch answering the research question. So the first question was about the impact of personalization, and by no means can I say something super generalizable here. This is a case study with Wayfair data, but the point I'm trying to make is to understand what is the impact of personalization in this particular case. So the way I estimate the causal impact of personalization using a large-scale field experiment. So together with Wayfair we took more

What I saw was also important is that usually when people talk about personalization, everyone remembers Chris Anderson's Long Tail literature where personalization should help consumers with niche preferences or very niche products. That's not what I observe empirically. Consumers who have very niche preferences, they usually do not benefit from personalization because they use very specific keywords because they know what they want and that's why they can find everything organically and very fast, versus personalization mostly benefits niche sellers, so think about small local sustainable wood or stuff like that, this type of sellers. So the takeaway here is that given the Wayfair experiment, what I can say is that the personalization could help consumers. It can help platform and smaller sellers on the platform. Again, I do not claim generalizability in the sense that any platform might have a different setting, but at the same time, from the regulatory perspective, it's important to know that there will be different outcomes for different platforms, right?

Now the next question is how will privacy restrictions that limit consumer tracking impact this benefits? Just to remind you, the problem occurs because usually consumers arrive for multiple sessions, 40% are recognized through logins, 60% are only recognized through online tracking technologies. For example, one of the privacy regulations that is discussed right now is a Safari first party cookie expiration policy where a Safari wants to set the first party cookie expiration date to seven days. So what that means is that if the consumer arrives after more than seven days, their cookies will be deleted. So just to give you an example, imagine there is a consumer who arrives for three sessions. After the first session, it took the consumer more than seven days to arrive back and then she arrived within a week, for example, for the third session.

What the policy would do is the following. The platform would no longer be able to recognize the consumer on the second session because the session originated more than seven days after the first session, right? So basically the cookies were already deleted and that's why the platform has to work with the fragmented data instead of having a full consumer history. To give you a scale of the problem, on average, 30% of consumers arrive from Safari browser, which means it's like millions of consumers who will not be recognized.

To understand the impact of this type of policies, because a lot of those are in the future, what I do is the following. So I take the original personalization algorithms that the platform uses and then I retrain the algorithm using the lower quality of data. So remember that I know the source of traffic and the browsers that consumers use when they arrive on the platform. So basically I split the data manually as if consumers who arrive from Safari browsers after more than seven days of inactivity were not recognized, and I basically retrain the whole algorithm with lower quality data, and then that generates the counterfactual set of rankings that would have been generated under lower quality data.

Now, the next question is how will consumers respond to this counterfactual set of rankings? For that, I obviously need to simulate that and I developed a session search model to understand that. The outcomes of interest in the model are the following. I'm interested how consumers choices will change as a result of counterfactual recommendations, what will be the outcome on the seller revenue and the platform revenue and profit. So in terms of the model, the model is as follows. So consumers arrive on Wayfair, they navigate to product ranking pages. They see only part of the page. So they have limited awareness and the reason is very simple because on your screen, you do not see all the products that are available. You have to scroll down to discover additional products.

Right away, in their awareness set, consumers observe the price rating and the image of the product. So those are observable product characteristics. They know their preferences towards this attributes. What consumers do not know is the product reviews and additional product details that they have to click on a product to reveal, and that's the origin of the search costs. Because some of the characteristics are observable and some are not, the true product utility is a priori unknown to the consumer. Follow

Hodgson and Newey's paper, I assume that this utility is a function that is a draw from a Gaussian process. The reason I do that is because it nicely allows me to incorporate consumer learning in the search model.

Because consumers do not know the true product utility, they start forming expectations, right? So as a consumer I might think, "What is my expected utility from clicking on a product that I already observe on the page versus my clicking cost?" Or I could go down and discover additional products, so basically scroll down and discover additional products if my expected utility from doing that is sufficiently high, or I could leave the platform overall and just go to the outside option, or I could purchase the best products that I've observed so far. So these are the four actions that consumers are choosing from. Obviously, because the model example is complicated and it is really hard to solve in closed form, so I assume a near

for the consumers who are more price responsive, versus consumers who are less price responsive,

who arrived from some IP address and they were looking for particular products on the website, then the next time someone else is arriving from the same IP address and they're looking for similar items, the chances are it's the same or very similar consumer, and then we can basically start showing personalization in a probabilistic way.

What I do next is I evaluate how would probabilistic recognition perform on the market using a structural model. Basically, the blue lines or the blue bars are still the previous Safari results that I showed you before and the green bars are the results from the probabilistic identity recognition. So what I find is that on average, all the results are going to be better under probabilistic recognition because now the platform does not know exactly who the consumer is, but they can probabilistically guess and still show personalized set of outcomes. So in terms of the key takeaways, what I learned from the project, and I hope I conveyed it here too, is that personalization is not necessarily bad. It can benefit consumers, sellers and the platform. Privacy regulation primarily is hurting smaller sellers and price-responsive consumers, and platforms can partially mitigate this losses if they use probabilistic identity recognition.

So in general, probabilistic identity recognition is sort of illegal in Europe because they don't want to profile people and there is a risk of profiling people incorrectly even when you use this type of algorithms, but I think it is one way to go where the platform still does not know who the consumer is, but they can probabilistically guess and still show personalized set of outcomes. Personally, I think it is way better than having consumers do biometric login because if you... I don't know if you've noticed, but some of the platforms, they actually introduce the biometric login option where you can log into the platform website using face or fingerprint ID, which I think is a different set of privacy violation versus this type of machine learning algorithms could help platforms operate under privacy regulation. Thank you so much. I really appreciate any feedback and thank you, Ginger, for discussing the paper. I'll click through the appendix.

Malika Korganbekova:

I will click through the appendix, sorry. Okay.

Ben:

And now as you may have guessed, we have Ginger Jin from the University of Maryland to discuss.

Ginger Jin:

Thank you, Ben. Thanks Pinar and other conference organizers for inviting me to discuss this interesting paper. I don't think I have any direct interest for or against Wayfair, but per Aviv's morning advice, when you're in doubt, you should disclose. And so I disclose that I provide consulting services to Amazon and Alibaba. And as far as I can tell, both marketplaces on Amazon and Alibaba has offered something, probably what we'll call furniture. I haven't checked whether they have exact, the same blue velvet dialing chairs. There are many antitrust experts in this room, so I will leave it to your judgment on whether this is relevant or not. But more relatedly, I'm very fortunate to have opportunity to work at FTC as the Director of BE about eight years ago. And that experience exposed me a lot to the privacy issues, to the data use issues. So I'm really glad young researchers like Malika has really picked up those issues seriously and writing such a fantastic paper.

I would like to encourage all of you, especially the students and young researchers in this room to really start thinking about the fascinating research area like this. So back to the paper, Malika already did a terrific job summarizing her findings, so I will be very brief here. You can see that she's addressing very important research questions, not only on whether privacy restriction would affect the platform, the

consumers, and the sellers, but also how the platform and consumers are going to respond to that changes in privacy restriction. I think that equilibrium view is really, really important. And she has extremely rich data. It's not only who purchased what, but also to what extent they click to scroll, to tap, or hover, zoom, return and repeat customers. So this is really, really extremely rich data. She also have a lot of experimental variations, the randomized experiment that shut down the data available for personalization, but also the price experiment and many, many other experiments that allow her to really customize her model for the Wayfair data.

I really envy that kind of data access. And you can see that she-~~Dap~~tm modeling and analysis, she

consumer search model from two samples. So one sample is for those who come into the store from first-party cookies, and the second sample is for those who come in from the third-party cookies. Okay? They do those estimations separately.

However, in reality, I would imagine that the same customer could appear in both samples, right? Sometimes, I use the first-party to get in Wayfair, sometimes I look at Weather.com and the third-party cookie there track me and show me the ads of Wayfair for the blue chairs and I click into it, right. And this could have, I'm sort of trying to think what kind of implication this could have for your estimation. I can think of two. One is the selection effect. Those who use both first-party cookie and third-party cookie are maybe those who are eager to buy and less sensitive to price. And maybe that's kind of why you find different price sensitivity in your sample. Another could be the ripple effect in your counterfactuals. If you block the first-party cookie in one counterfactual, it may hurt the Wayfair's ability to use their third-party cookies on the same consumer. So I don't know whether there could be ripple effect in your estimation.

My third comment is about this probabilistic identity recognition counterfactual. That's kind of the counterfactual algorithms you think Wayfair could use to get around of the privacy restriction. I think the motivation on this could be better because I can see the motivation of the other counterfactuals very tied to Safari's policy or Chrome's policy, but this one seems like you've just kind of hypothetically assumed that you cannot track the same users across different devices. I'm not aware of any real policy that try to limit of that somehow about (you) (3) (4) (5) (6) (7) (8) (9) (10) (11) (12) (13) (14) (15) (16) (17) (18) (19) (20) (21) (22) (23) (24) (25) (26) (27) (28) (29) (30) (31) (32) (33) (34) (35) (36) (37) (38) (39) (40) (41) (42) (43) (44) (45) (46) (47) (48) (49) (50) (51) (52) (53) (54) (55) (56) (57) (58) (59) (60) (61) (62) (63) (64) (65) (66) (67) (68) (69) (70) (71) (72) (73) (74) (75) (76) (77) (78) (79) (80) (81) (82) (83) (84) (85) (86) (87) (88) (89) (90) (91) (92) (93) (94) (95) (96) (97) (98) (99) (100) iPhone to a Samsung tablet, right? Maybe a better motivation on that would help readers to understand that. And some of the results seems to suggest that this algorithm that's not as good as the full level of tracking somehow can benefit some consumers. So maybe, could this be some of the things I was talking about? Maybe flesh out that a little bit would be good. Yeah.

Maybe this could be linked to other privacy restrictions, especially how this algorithm could address, say, Safari's limit on first-party cookie or Chrome's limit on third-party cookie. Those counterfactuals

Speaker 23:

Now, we have time for some questions.

Malika Korganbekova:

I think there is a question here. There.

Speaker 24:

I'll start. On your right.

Malika Korganbekova:

Oh, there. Yeah.

Speaker 24:

This side first. Tremendous paper. It's so interesting to see the ways that privacy in practice personalization had such big benefits. And I'm wondering, probably not in this paper because it's such a

Speaker 23:

Any other questions? All right.

MalikaKorganbekova:

Okay.

Speaker 23:

I think we can move on.

Malika Korganbekova:

Thank you so much. Thank you.

Ben:

And next we have Evan Starr from the University of Maryland presenting Clause and Effect: Theory and Field Experimental Evidence ~~Non~~-Compete Clauses.

Evan Starr:

Okay. Thank you so much to the organizers for having us here, and thanks to you all for sticking around. This is the last presentation before snacks. So I hope by the end of this you will say to yourself, I'm glad I stayed. Okay. And if not, then at least you'll get snacks. All right. So this is joint with Bo Cowgill, who's sitting here in the middle, and Brandon Freiburg from Columbia. And let me first disclose, this is a field experiment. We got some funding here from the Russell Sage Foundation, from the Smith Richardson Foundation, the Institute for Humane Studies, Columbia Center for Political Economy, and the Copper Foundation. And a full disclosure, I've also been retained as an expert witness in several labor market competition issues over the last few years. None of those have any interest in this particular paper. Okay. So first, I probably don't need to explain this to anybody in this room, but we should start with what a noncompete clause is.

It's a term in an employment contract that prohibits a worker from starting or joining a competitor firm within a particular timeframe after they leave, usually a year or two and within a geographic boundary. Okay. So here's a ~~non~~compete from Amazon where they're prohibiting a worker for 18 months after they leave from engaging in or supporting the development, manufacture, marketing, or sale of any product or service that competes or is intended to compete with any product or service sold, offered, or otherwise provided by Amazon. And so these are controversial restrictions. The folks in this room of course know that the FTC took a position on this in the last few years, and it's been a controversial position that has been playing out in the courts recently. And so I want to take for a moment the position of the case for ~~non~~competes.

Over the last few years, especially with the FTC comments that we've had, the case ~~for~~competes has been made relatively clear. So what are the main tenets? The first argument that you often hear is that non-

with a simple model of optimal wage setting and ~~no~~ competes. And the question we're going to ask

The field experiment looks like this. Okay. So in phase one what's going to happen is we're going to work with what we're going to call Firm A and we're going to randomize ~~compete~~ and wages in job offers to sent to 14,000 individuals, and then we're going to hire workers who take the jobs and then we're going to give them some secret information to do a task. Okay. After they're done, then we're going to wait a little bit and we're going to work with a second company. And that second company, who's a competitor, is going to try to hire the workers who were previously hired in phase one. And we're look at who joins the second company and then who shares the secrets. Okay.

In the third phase, we're going to have workers who are randomly in there, they're going to be in their non-compete for three to four months and then we're going to randomly start releasing them early. Okay? And then we're going to do a follow ~~up~~ on them about a year later. Okay. So here's kind of the punchline results. What we're going to find is that ~~non~~compete reduce earnings and mobility, but they don't protect secrets any more than ~~non~~disclosure agreements. Okay. So this is going to be our punchline finding. We're going to find mechanisms that relate to a lack of contract reading and that a reminder of worker obligations is key for these results. Okay. One other kind of fun result that we didn't expect is that when you make a ~~non~~compete very salient to workers, that it selects workers who are more willing to break them. Okay. So if I put a ~~non~~compete in your face, the workers who are willing to sign that are less likely to abide by it. Okay. All right, so that's where we're going.

So let me talk about the theory here briefly. Okay. So we have a very simple model of contracting under uncertainty. We've got two players. There's a firm and a mass of workers. Okay. The workers are going to have some private distaste for the job, D . It's going to have some distribution, F . And the non compete is going to reduce the net present value of future earnings by some amount, K , and there's

are the workers who are then wrong about enforcement when they assumed it wasn't going to be enforced and then it actually is. Okay. The other thing you can do is you can add uncertainty about the presence of a non-compete agreement. All right. So if workers don't even read their contract altogether, then there's no compensating differential for the non-compete, and then workers who have unknowingly signed a non-compete become worse off later because they suffer this penalty, K. Okay. And then behavioral factors are going to make all of this worse.

Okay. So that's the model. So let me turn to the experiment, which is really the bulk of all of this. The population that we're studying is contract HR recruiters. Okay. And so let me break out why we're studying this population very briefly. Let me talk about the contract part work, worker part first. Okay. So contract workers are sort of interesting because they should be especially unwilling to give up their freedom to work because their ability to make a living depends on their ability to work with multiple employers. Okay. But they mostly don't interact with each other, which makes SUTVA violations a little bit less likely. Okay. They're of independent business and policy interest. If you go to LegalZoom, you'll see articles like how to use a non-compete when you work with independent contractors. In terms of policy, here's the... The New York City Council proposed a law to ban non-competes for freelancers.

They say, covenants not to compete are increasingly common in contracts between hiring parties and freelance workers. And they have a bill that would prohibit non-competes for those workers. Okay. What about HR? Why are HR workers important? Well, HR recruiters are a growing part of the U.S. labor force. They have access to very valuable information because they know who works where, they know who's potentially mobile. They sign non-competes at a relatively common rate, slightly more than average, about 20 to 30%. Maybe most importantly though, they're aware of non-competes and the harm they can cause because they're engaged in hiring workers. Okay? They're also may be experienced in bargaining. All right. And so, what do we think here? We think that by setting this population in any field experiment, you're limited in terms of generalizability. Here we think that, we think contract HR recruiters should probably be more averse to non-competes and less likely to be harmed by them than at least the typical worker.

Okay. Let me tell you about the sampling frame here. What we did is, we worked with Firm A to identify about 30,000 HR recruiters on a platform. We took a stratified random sample of those 30,000, and then we're going to inverse probability weight our sample to reflect our 30K population. And here's what our sample looks like. They have an asking rate of about 50 bucks an hour. Most of them have experiences in recruitment and finance. The one thing I'll highlight is that we do a variation in the states where these workers live. And so the bottom row there, 70% of workers live in states where their non-competes are potentially enforceable, but in for 30% of the workers, their non-competes would be entirely unenforceable in their states. Okay. So we'll exploit that later. Okay, here's the experimental manipulations. So when it comes to the non-compete, the control group is we have no non-compete. Workers have an eight-page employment contract and the signature is required at the end.

Our first manipulation is that we include what we're calling a hidden-compete. We're using the word hidden here because we're going to contrast that with a salient-compete. You should think of this as the normal condition. Okay. This is a normal-compete. It's identical to the no non-compete, except until you get to page seven where there's a small little paragraph that includes the non-compete. Okay. And then we have the salient-compete, and the salient non-compete is mentioned in the job offer. It's right up front, it's on the very first page of the contract, and it requires a separate signature. Okay. So you can't miss it. Okay. All contracts also had disclosure agreement. They used Florida law where these firms had offices. And all signatures were checked. Important details because we are the ones working with these individuals, we were able to record how many milliseconds they spent on every page of the contract. And importantly, there was an option at the bottom of the contract to click go to the end and sign. And if you did that before you hit page seven in the hidden condition,

then you wouldn't know that you have a non-compete agreement. Okay. So here's what the non-compete agreement looks like. This is the salient page one. It says, during the paid engagement with the company and for a period of six months after, the recipient will not directly or indirectly engage in any business that competes with the company, including but not limited to, business engaged in finance,

When it comes to this, we can say this is about 8% drop. And actually because we randomized wages, I can tell you this is equivalent to about a \$6 an hour difference. Okay,. In the salient group, the salient group, about 15% of them opened the contract, 13% of them signed it, and 10.2% of them completed the task. And we can reject that difference from the no-~~compete~~

Evan Starr:

We did the task and we can reject that difference from the no-~~compete~~ group. That corresponds about a 15% drop in task completion rates, and that's the equivalent of about \$13 an hour in terms of the initial job offer. You might want to know who's driving this. We have a lot of heterogeneity in the paper. I don't have time to get into it, but I'll tell you what's driving this first stage result is, it's women are the ones who are more sensitive to a salient-~~compete~~. It turns out men sign at basically, similar rates regardless of what's put in front of them. And people who have high asking rates are also deterred. All right. What about reading the contract? So here's the distribution of the number of seconds spent on the non-~~compete~~ page, conditional on opening the contract and getting to that page. So the salient group is in the gray bars, and you can see that there's a huge mass at 60, which means that most people spend over a minute on the first page of the contract. And contrast, the hidden group, there's a huge mass on the left, and the people who spent zero seconds are what we call the skippers. And you can see these are about a third of the hidden. They skipped the-~~compete~~ altogether. On average, about 75% of workers spend less than 10 seconds on the hidden page seven. Okay?

firm B is going to come in here as a competitor to firm A. After the engagement with firm A ends, they're going to make offers to every worker that was hired by firm A, and they're going to randomize the wages between \$27 and \$62 an hour. The task for firm B is that firm B is a similar finance company, and what they need is they're sourcing leads for an opening. So they're trying to find leads. Everybody who is in this experiment could give them leads. All they have to do is share the resumes that firm A shared with them. So this is going to be our measure of secret sharing.

So what we're going to look at is whether workers accept or not firm B's offer. We're going to see if they share resumes. And then part way through this engagement, firm A is going to send a reminder message

enforceability of the noncompete. So if you're in a state that could possibly enforce your noncompete, this is the pattern we observe. It's the same as before.

Now on the left I'm going to show you the same graph, but these are workers in states where their noncompete would not hold up in court at all. And here's what we find. You can see it's exactly the same. All right, what about sharing firm secrets? So overall, before the reminder happens, we find about 4% of workers share secrets. We can't reject any differences with noncompetes. They share at similar rates. After the reminder, we find the same thing is true. After the reminder, people share less. But even though the noncompetes dissuaded mobility, it didn't add any extra protection in terms of protecting secrets.

So Phase 3. Phase 3, what we did is we are going to randomly release workers early from their noncompete and give them another chance for follow-up work. So let me just show you the result here. So what's happening in this case is everyone's getting sent a job ad and given the opportunity to apply, but some of them have not yet been released from their noncompete agreement and others of them have.

So on the left here, I'm going to look at these are workers who don't have a noncompete, and some of them have gotten... So everybody gets a reminder here, and the people who don't have a noncompete get reminded about their NDA. People with a noncompete get reminded about their NDA and released from their noncompete. So if you have no noncompete, it doesn't matter whether we send you a message, a reminder at this point. If you do have a noncompete but you haven't been released yet, we find you're far less likely to apply for this job. But once you've been released from the job, you partially respond and move upward. So this suggests that again, the noncompete was holding workers back from taking this over job.

Lastly, what we did is we waited a year. It's very painful, but we did. We waited a year, and then we scraped their earnings on this platform. And we were curious, how does the noncompete agreement... If I randomly offered you a noncompete agreement versus the same job with no noncompete, or even no job at all, how does that affect your earnings? And the punchline finding here is that if I offer you a job with a noncompete agreement, we find that it reduces your earnings by about 4.5%. This effect is driven almost entirely by the hidden group. They have an effect size that's about twice as big as the salient group, and for salient, we can't reject a zero here. Let me just jump to the last column here. For the skippers, these are the people who skipped page seven, they drive an enormous part of this earnings loss. We find that they have earnings losses about 21% if you skip the noncompete entirely. Again, this is on the platform.

So let me talk about generalizability, and then I'll wrap up here. So what do we think we might generalize to? So HR contractors we think are probably more sensitive to noncompetes than the average worker. What we think here is we've studied secret sharing, but only with regards to things like cli the r

competitor company, and then they're following up after a year. They're also collaborating in the context which is not standard. This is ~~not~~ compete.

As someone who runs field experiments, I was so excited to see this being implemented. So really hats off just in terms of executing this. Also, I think important insights. As Evan mentioned, a lot of the literature focuses on natural experiments that are created by state laws, and this was a really cool way of leveraging digital footprints to highlight the fact that people don't read that much and people might not be aware. And that I think really adds to the conversation.

Also overall, I buy the findings. For example, the reminder intervention, which prevents job mobility, I think it works on exactly the workers you think it should work on. So overall convincing ~~with~~.

I've gone through two, two and a half versions of this paper, and so I encourage all of you to read it. But of course as a discussant, I'm trying to get the authors to tie up the loose ends. So just in terms of the context, as a reminder, the context is a ~~coffee~~ one-hour job on the platform. So when we see minimal wage bargaining and explicit sharing of client lists, how much can we extrapolate? So this is a downside of field experiments in general, but I think a deeper conversation around this could be helpful. Of course, when companies' names are anonymized, it allows you to partner with them, but then there's a gap in extrapolation to other contexts that we might care about.

Another thing would be to actually leverage some of the heterogeneity, maybe by experience to shed light on the role of beliefs and information that might be driving the results. I know that you've done some survey stuff, which I really liked. So now that the experiment is over, another wave could be to just survey people to understand what their beliefs and information is, including what they might be doing

Speaker 27:

Thanks. This is just a super cool paper. One suggestion I had, if you haven't done it in the paper itself, is to look at your standard errors and try to talk about how precise your nulls are. What effect sizes can you rule out? I think that could help understand exactly how to interpret some of the null results here.

Evan Starr:

Yeah, that's a great point. Yeah, we could work on that for sure.

Speaker 28:

Could you say more about whether these noncompetes are usually used in this industry, and how that compares to others? Would people be expecting that this would be in their contract or not?

Evan Starr:

So I mean the finance industry actually is really... The finance industry in New York blocked a non compete ban from coming in the state of New York, which was passed by the legislature, and then they lobbied the governor to keep noncompetes. I mean, our sense is that noncompetes are actually quite common in the finance industry. And with HR recruiters, the companies that we're working with have used them with HR recruiters in the past. So this is something that's happening. It's happened with freelancers, it happens with hairstylists who are independent contractors, it happens with yoga instructors. All sorts of independent contractors have these sorts of agreements. It's relatively common in our context. I mean, at least in line with national estimates.

Speaker 8:

Hey, so different US states take often widely differing approaches to noncompetes. They have different levels of restrictions and so on. California for example, bans them completely, and you have Silicon Valley over there. But then I wonder, you also have states like Wisconsin, for example, that have a smaller but still thriving tech sector where noncompetes exist. I think Massachusetts too. To what extent, I wonder is that differentiation possibly a good thing in terms of providing a market for companies to go and set up and create jobs otherwise wouldn't be there because they really want non competes that badly?

Evan Starr:

That's a great question. Part of a bigger discussion about the role of federalism and navigating cross state issues. I'll just mention, of course, variation is g 1 0 0 >c, ofr9havn7/MCIslfe2(is so) w6(te es can)- rule. P

Evan Starr:

All right. Thank you everybody.

Speaker 26:

We'll wrap it up.

Sam Kleiner:

Okay, everyone just wanted to thank everybody for coming. That concludes the day's events. Please join us tomorrow at 9 A.M. for the remainder of the conference. ~~There's~~ There's going to be a reception outside with food, and I wanted to also just thank the Tobin Center at Yale for providing all the food. Hope to see you both outside right now and tomorrow morning.