

16th Annual Microeconomics Conference Day 2

All right,

[>>] we are going to get started here all right, good morning everyone, welcome back, I hope You all had a pleasant evening last night. We are going to jump right into things, we are going to start today with Mike Sinkinson. Mike Sinkinson is an associate professor at the economics Yale school of management and specializes in industrial management, he recently completed a year of service focusing on competition policy and tech policy, he will be providing us with a few opening remarks, and then we will follow with his keynote address, advances in testing for the nature of competition.

Sinkinson] Hi everyone, thanks for being here on a Friday morning after hopefully a fun evening. So I wanted to begin actually by apologizing that I'm not Steve Berry because he's supposed to be here to give a welcome on behalf of the Yale token center economic policy, unfortunately, he's in New Haven trying to find a new president for Yale university, it's like an important thing to be doing. So he asked me to stop in and get a few remarks on behalf of the Tobin center, so Tobin center's mission is to support research that has an impact on society, that's kind of the goal of the mission of the organization, so the center is incredibly proud to sponsor this event in particular.

Because this event brings together top academics as well as government practitioners people from agencies. And so forth to spread that research that's happening throughout the government and to make it useful. So Heidi yesterday on her keynote mentioned, she was talking about the topic of high skilled immigration and how we didn't have the research that we needed to make the policy case that we wanted to make, and that really resonated with me, because I just served on the council of economic advisors and one of my issues was have you heard the president talk about the war on junk these I was one of the soldiers in that war and let me tell you, it was really important that we could point to high quality research on the topic.

So we could point to Allison and Allison 2009 on price obfuscation, we could point to Gavin Liebson about shrouding of attributes kind of we could point to cautious times work with David Muir and Ariana Mitrona and drip pricing, we can point to that research and say look, this is real harms. This is not some fool's errand that we are on. Come we actually care about addressing junk fees fees. So the Tobin center is very happy to be sponsoring this event. I think a lot of the papers we saw yesterday think you all agree with me that definitely have impact in terms of regulators agencies and so forth, so the opportunity to sprint that research throughout the government is a really great one so yes, the Tobin center is very happy to sponsor this event, and alright my Steve and my Steve Berry impersonation.

So I'm going to jump right into my keynote advances in testing in testing for nature of competition comment I want to clarify this talk is mostly going to be based on a research agenda with a couple of co authors Matthew Backus and Chris Comlin at NYU. This work goes back to the 80s basically, they were thinking about cartels and collusion could you detect it. There was subsequent work that used menu approach or the work for example, saying is this perfect competition, as this monopoly, is this collusion? Can we compare different models of conduct? More recently, there's been a revival of this idea of estimating an internal parameter, so for example Matt Winebrook's work in fear or the Crawford paper and television.

They were trying to estimate what is the weight the one for places on another firm's profits, how much are they in terminalizing another term's process, very related to what we are doing some parallel work by Duarte McNulty solstin and Sullivan, they are thinking about different statistical testing frameworks, and which ones performed the best, they do a number of Monte Carlo tests and they basically find that the rivers and long framework which we will be using and I will explain in a little bit does perform very well compared to, let's say, a Cox framework. And then towards the end of the talk, I am going to talk about some other framework that we've used.

And so we obviously love that, what's nice is we have a kind of spoke context that we are going to develop our test and I will tell you more on the next slide, but the test has broader applicability, so for example, the stark and woolman paper is thinking about collusion in generic pharmaceuticals, the resilient papers think about how wages are set in a labor market, the car older wing and Kim paper it's about collusion in rent,, I was looking into historical context into conduct premier so the one thing I just want to say, though, is conduct, even testable, is something that could be identified and very hail 2014 they have this very nice sentence that says you know, you could use exclusive restrictions and supply relationship to test for conduct, so a lot of this project was trying to figure out the best way to do this kind of work OK, so again, absent additional restrictions.

I can't just look at data on price quantity and decide whether or not say collusion is taking place. The problem is there's a lot of different models that could give you the same outcomes. Could look a lot like collusion rate but those have very different implications. So in our work, we a lot of different ways, you will see what that means, then second, we are actually going to use the berry Hale idea to use exclusion restrictions and supply equation to differentiate between different models of conduct. If you have taken graduate IO you know this picture, this is differently canon for IO economists right, this is Brosnan 87 giving intuition to how you might detect collusion.

So you see on the x axis are products made by different firms and products to injury in particular are very close in product space, they are very close substitutes. And then the left picture there prices high above the marginal cost curve and that would be in the collusive world where so the right picture in the competitive world those two products and say, are so close substitutes would have very low markups they would be very close to the marginal cost curve, so this was the intuition that he wanted to build. He said these models have different implications for behavior, and then his work.

He basically was the one that coined the intuition of you need rotations of demand to identify conduct. I just want to highlight another few papers in this literature, Sophia has this next paper in 2007 where she compares a whole bunch of different models in a grocery store context, in particular she's considered a double marginalization which she's doing as basically estimating the supply relationship with G&M and comparing the GMM functions across the different models and saying what fits the best, she then takes that and puts it in a Cox framework. So this is kind of her main testing table, the challenge with a Cox framework is you can get contradictions from a some particular, in her test come up.

You're a null hypothesis is that model 4 is the truth, you reject model 5, but if your hypothesis is that #5 is the truth. And you reject model 4. So that's not exactly the great property for a So test. as I said, we'll be using your rivers in Wang framework instead. And as I said, I'll show you how to set that up. So one more paper from the literature because this is the one that we kind of. of build off It gives the following example that I want to walk you guys through, so the picture on the left, you have an equilibrium outcome E that you observe in the data, quantity and There's two price.

models that you can consider in which case you would infer that the model cost is MC_0 , alternative model is that this is a monopoly firm or with market power in which case you would infer that the marginal cost. Cost is MC_1 . Now, suppose you had some insert mints that rotated marginal value, however, they are excluded from supply, they are excluded from cost, so that's the right picture, we are going to rotate revenue. revenue what you get is the dashed lines, and you have a new outcome ET prime, now note, though if you think the model is perfect competition coming you would now think that the marginal cost is MC_0T prime, it has moved down from MC_0T , whereas if the true model is monopoly, the marginal cost is unchanged, well as we just said, this is an instrument that is excluded from supply, it can't affect cost, so right now we can reject perfect competition using this instrument.

That's the intuition behind the approach that we are going to use, we are going to have instruments that are excluded from supply that are actually rotating demand. OK, so as I said, the meat of today's talk is going to be about a paper that I have with Matt Backus and Chris Conlon where we are testing the common ownership hypothesis in ready to eat cereal. So I'm not going to cover the hypothesis come I guarantee everyone in this room has seen a paper that talks about the common ownership hypothesis and I'm not going to go through the math of it or anything like that, but the basic idea is the following, if my largest investor is my competitors largest investor maybe we won't compete as intensely, since we've all been told for many years to buy the market to buy index funds that are diversified, there's a big asset management companies are the largest investors in every publicly traded company for part. So the idea is if my largest investors are the same as my competitors largest investors we won't This is the kappa FG on firm G, so these profit weights are based on overlapping ownership and some other assumptions about how the firm irates and preference investments, previous paper we actually showed that the profit waves provided by profit ownership have grown significantly over time because again, this indexing strategy that we've all moved towards, actually we document that the S&P 500 take two firms at random, on average, they will have profit waves of .7 on each other to an IO economist that's a scary number because when we think about a profit weight and another firm, we think 0 is competitive and one is merger right?

So that's a big number. So they really imply an intermediate model of competition between our standard owned maximization and full collusion. So, our goal is to come up with a test that's able to distinguish between models of competition that are close so, not to discredit any of the earlier work, but detailing perfect competition apart from collusion those are far apart, they look very different. The challenge we face is that this can look close. So we need to have a powerful test to distinguish it from other models. And we are going to do this in ready to eat cereal cue the jokes that IO economists only knew for economies.

But we've picked cereal because we think it's a nice laboratory for this, I will show you this in a slide Kellogg's one of the biggest players has its largest investors the Kellogg's foundation which is undiversified so it has very stark predictions under common ownership that are different for Kellogg's and the other firms. So we're going to have some very nice variation in that measure. OK commas I'm not going to go through the full notation the full demand model, but you should expect that we have differentiated product set up markets are indexed by T, by J, we're going to observe product characteristics for all of our goods, as well as market shares and prices, in addition, we are actually going to have an augmented model where we observe consumer demographics, so we will have as you might

imagine what's important for cereal is, are there children in your household, so we observe that as well as income and we are able to observe those in our model as OK, so well.

what does the common owner hypothesis mean in this world. We are just going to generalize what's called the ownership matrix, which is usually 0s and ones for which products a firm owns to account for the profit weights implied by common ownership. So what our firm is doing inside their objective function is the kappa weighted prefix of competing products so in some sense, I care kappa about your profits. So I care about diversion to your products now. So instead of zeros and ones in our ownership metrics we're going to have the profit weights from the ownership world representing how much a firm F cares about the firm G, when you go to the data we're going to do this for bridge scanning data to get the demographics we are going to have, like 10 years of data, we have plenty of data.

One two things that the estimated model allows for very flexible patterns between products. So our test is going to come down to what do the markups look like in different models, so we need to get those diversion ratios right, because again Kellogg's is going to behave differently than general mills. We estimate demand we treated is known basically. And then we can enter back up So what's markets. the challenge that we need to overcome? The model for true Markov will satisfy for the supply equation given the screen near, so on the left hand side, we have price minus market, otherwise marginalized cost is on the left hand side is equal to some supply equation $H_{sub S}$ plus an omega term and the conditional moment restriction is that is equal to 0.

That's how you get this. The goal is say to we have two competing models of two competing markups and we want to see which one seems to fit better. But there are a lot of challenges in terms of implementation, the one, the test is going to depend on how we choose our unconditional moment restrictions come in reality. What do we do? We recovered the omegas, so how you choose a is going to matter. In addition, it may depend on how we choose the $H_{sub M}$ function, so there's a hand that says aren't you really doing a joint test of your specification and a conduct model, let me just say we will overcome that.

And we will not worry about that, in the 2017 for what she does if she tries log linear exponential and X beta, and they all quantitatively agree. So then she moves on. But that's a researcher choice. That's unfortunate in this And then again, setting. your choice of hypothesis may affect the weighting matrix that you're waiting for your test and that's the test itself, the Holland palace critique critique. If I want to motivate how we kind of solve this coming from of all paces missed specification, so index the true model of conduct by zero, and the first equation on the slide price minus markups 0 we know that equation is true, it holds, it's correct, so to motivate a useful test, we ask what happens if we estimate supply with the wrong conduct model, suppose I put price minus markup one on the left, side, then I estimate that, I will recover an omega one term, what am I actually recovering, what is omega 1?

You can it break down through a little algebra to see it's actually the true omega 0 plus the true markup minus the wrong markup or the difference in those two markup models. So what when happens you miss specify your conduct model? We're introducing an omitted available at the omegas he recovered correlated with the difference in markups, because you miss specified, so that's going to be the basis of our test, so this is our innovation. This is kind of what we came up with, so the model is going to be given by price minus markup equal supply minus omega, and we have some unconditioned moment restriction where we interact with the megas with functional instruments,, what we suggest doing is choosing the function A&Z which is the expectation of the difference in the Markov's projected on instruments instruments.

Many advantages, all of a sudden, we produce potentially a many moment problem the a function to a single scalar moment, it's a number. no need for There's a waiting matrix we don't need to worry about that at Second all. testing is reduced to Two markers we are looking on instruments that we are recovering the omegas and as we can show you we can do that very flexibly, we are going to do that on random force, so there's going to be no functional form products at all, we show in the paper a couple ways that you can list this if you're comparing GMM functions.

This approach is going to maximize the difference between 2 mmm objective functions conditional on any matrix. So we think this is a good approach. So let me give you quick overview of a how we do this and we are working in the rivers of long framework, which is a non nested model framework,, we assume demand is known. So Eight eight O O 1 2 are the markups for model 1 and model 2, and for each model, we are going to complete a function function. This is going to be Eric you've come and what we're going to do is we're going to interact with our omegas with that expectation that the difference in markups projected on markups, we are looking for violations, we are looking for that that correlation tells us that there is a problem the null hypothesis in the rivers and bong framework is that both models are equally close to the So we truth are testing 2 models OK?

the nil says they are closed, the alternative has set 1 model significantly closer to the truth than the other, what rivers in Wong show very helpfully is that this test statistic right here. Once you scale it properly has a normal 01 distribution, standard normal, so when I show task statistics. Statistics you should be thinking plus and -1.96 from like first year stats right? They are critical the values of this test that we will tell you that one model is closer to Getting the that scaling factor truth. right. Of course is really difficult come there's no easy way to do it, so we bootstrap our application application.

If so I said we were going to use instruments to do this. So we need variables which affect demand, which rotate marginal revenue, but are excluded from supply, and aren't you glad no I didn't go through the full demand specification and go through all of But this this is the first is notation? demand and the 2nd is supply in our model, so, what are some variables that affect demand. But not supply. So in our notation. These are variables that are literally only in demand not in the supply equation, we tried things like product recalls, and we found actually that they are just very weak, there's like 3 product recalls in our entire sample.

So we didn't get much action from those variables, the second thing that you can think of using is the YS which are the demographics, I think this is very nice because clearly, if you have wealthier consumers. That's going to affect marginal revenue, but it doesn't cost me higher marginal cost to to sell people cereal who are wealthier, so excluded from supply, but relevant to marginal The third thing we use revenue. are the BLP instruments. This always takes advantage of what else is in the choice that, clearly that's going to rotate marginal revenue, but again, it should not affect my actual marginal cost.

And then really nice one a that we like is the cost of other goods, so imagine I'm Raissa krispies, and there's a cost shock to oats, that will affect the price of oat brand right which affects My marginal revenue curve. But does not affect my cost directly, so it's excluded from my cost. But relevant to demand. So these are different types the of instruments that we're going to use in our application, we don't use the under demand shocks XI which actually affect demand. But not supply. But we have to assume that they are uncoordinated with the omegas to use them and so we don't want to make that assumption, we also those don't use the actual cap as because the finance people told us not to, and so we are not using those.

All right. So here's algorithm, we estimate that a function under models one and to get the omega residuals, and as I said, that H sub's function we're just going to use a random one we don't care we're not picking a functional one our hands are tied, then we estimate the first stage regression for delta ETA or the difference in markups between the two models, again we are actually going to use a random force for this too, and for each model compute the value of that criterion function Q, which is pretty straightforward, you're just multiplying the omegas times the estimated random force and then keep doing that for bootstrap samples to get the standard error of the difference, and then complete your test statistic, so actually once you've estimated demand that's the hard part doing the test is not the hard part here, it's actually pretty straightforward.

So not everything is testable, we are pretty upfront about this, the difference in markups between two models cannot be explained by instruments, this does not work here. What do I have in mind there, imagine we are testing markups the differences in markups is constant, so there are certain comparisons that we just can't do, we can't distinguish between them. I also want to point out the flexible demand model is very important to get cross sectional markup, it's where you actually need to have a good sense of what the diversion ratios look like between two different products that the markups are the same under all your different models a plain logic is not going to be very useful demand model model for our test. are doing We everything as flexibly as possible. to the main data that we Welcome are using it is the data from 2017, we consolidate ones to the DMA chain week level. So a market is going to be I'm in DC. So giant in DC for a week, right? So we do that, because these pains that the DMA chain level the price is typically the same, and we think that these chains tend to locate in similar areas, so the demographics are kind of the same at all We of cheat keep their only the stores. ones that do this and keep observations from a panelist data.

So we can actually observe 500 families that have shopped at this store, so we have a good measure of demographics. In terms products we consolidate from of the UC level, we work to scan a data Camino with a detailed, we go up to the brand, so honey nut Cheerios as a product, but that's different from other multigrain Cheerios. So we don't go the way up to Cheerios we all go one step down, it turns out there's a lot of products that have a lot of variety, citing special K has 21 different varieties in our sample, and not all of them are that healthy.

See how come we go up to that level of product former price is really just going to be revenue divided by servings, and we maintain the fiction that households purchase servants, we get the market size by looking at purchases of milk and egg servings and just basically scale up from there. So show you why cereal is such let a great setting, these me are the Four big players general mills Kellogg's quaker oats and post, and what I'm plotting are the prophet wade's place on themselves and other firms, so the flat line at one is always myself, I always place a profit weight of 1 on myself, if you look at the top right first you see Kellogg's is very different from the other, this is what I mentioned earlier, they have a large undiversified owner called the Kellogg's foundation and said they don't care about profits of other firms they have very low profit weights.

Weights on their competitors can contrast that would say general mills in the top left, they place profit weights of .5 to .75 of their competitors, so they should be taking in account competitor profits when setting price under the common ownership hypothesis, quaker on the other hand sometimes has profit weighs greater than one on its competitors they would want to to tunnel their profits competitors if they could, and I will last mention posters because post is adventurous, as part of altria, spun off to craft, as

well by Ross and Karina, it was iPod, those are big financial market events that should have nothing to do with the product market come up.

But they generate that big change in those profit weights over time, so what we did is estimate demand. And then predict markups for each of these firms under three different assumptions, that the blue line reach firm is the distribution under own profit maximization of their product markup, the red line is monopoly, and the yellow line is common ownership, so as you see exactly as we predicted as an intermediate model between the others the common ownership falls in between, but again, differentially for different firms, if you look at Kellogg's the yellow and blue are very close, or if you look at quaker oats it's the yellow and red that are very close.

So very different implications of the Connor common ownership for these different firms, this was just a kind of exercise of how big is common ownership or how big do we think it is, what we just did with our demand model is assumed they were playing own profit maximization, and said OK, now let's simulate A merger and see how much would go up So the first six columns, by. we do pairwise mergers and if you just look the at the bottom price row at index industry you can see the two biggest firms merge the price of 5 1/2%, then if we go to the very end, you see that if we turn on common ownership, the campus CEO, if we turn actually that on prices would go up by Shaputi nishik on more by common than ownership. It that. would be big these firms have high profit ownership on each other. This was just to figure out is this common firm big in this setting. OK this is our main results take one. This is a lot of numbers. So I'm going to apologize. I am going to take care. Few minutes to go through this, they'll come every one of these numbers that you see in the table as a test statistic from our test, which I told you is a standard normal distribution a normal So again, the critical values of 01. the test are ± 1.96 or minus.

So what the three panels are is we are building up to our full specification basically, what you see in panel one, we're just throwing the instruments in and having a linear function for the supply equation, what you see is across different columns each column has a different set of instruments, what you see is actually, we don't get, there are great results, so by the way, I should have said a negative number favors own profit maximization and a positive number favors whatever model is in the row, we test common common ownership ownership as a moving average common ownership with a lag perfect competition, and monopoly.

So in the first column, we can reject common ownership models for as advisors statistically significant and we can reject monopoly but we actually fail to reject perfect competition in that model, demographics. We don't get much at all. We can reject perfect competition. But nothing else, the BLP can reject not common ownership again, common ownership can be very close to the profit ownership maximization model, they work surprisingly well because they are a combination the first three reject all of those models come. And then in channel 2, we do the first part of our trick and say the AFZ function is going to be the the expectation difference of of markups and instruments, we are still in a rigid world where we have linear H&G functions, and the results are similar or if anything a bit worse, and in panel three, we do our full test, what's kind of remarkable is in all of our common markups we use, we get pretty consistent results for our test statistic rejecting all of these models versus profit maximization, we find in favor of own profit maximization using our test over these other ones, but notice it really did take our approach where we said, where is the information that's going to be relevant, competition what's the difference in markups conditional on instruments and so by just focusing on that as a test, we can actually reject much more OK, every time we present strongly.

this someone raises their hand and says, what about an internalization parameter, what if only sum of common ownership is happening right comic there's a friction, what if half is happening, suppose we just put a little parameter τ in front of the profit weight forms problem. And now this is an extra parameter where if it's zero we are back to our own profit perfect maximization camera if it's won. They're doing common ownership. And if it's somewhere in between, it's some friction right? Again this is something we have added to the paper because people keep asking us to do it, we're actually eventually happy that we did because we got this next really cool plot out of doing it, so what I'm plotting here on in the x axis is the test statistic, on the Y axis, I'm increasing τ .

So I'm going from own profit to maximization to full product ownership, each of these lines is our test statistic under different types of instruments and the real big take away here is I can't reject τ of .1 or example .2, because those two models are so close to own profit maximization I can't tell them apart, but, I can actually reject a tower of .3, our model is actually able to reject the 30% or more of common ownership incentives are reflected in the data, so this was actually coming. We saw this is definitely good news because we were actually able to tell my shows that our test is actually very capable of distinguishing these models.

OK, so, I want to use the last time I talked to broaden the discussion a little bit, so what did we do in the paper, we developed this testing procedure based on the result for from Barry and Hale about exclusion this restriction, rainbows that rotate demand but are excluded from cost, but keep in mind, our approach can test any pair of models such that they predict different markups. If me your favorite model of common you ownership we can test it come give if you give me your favorite model of collusion or behavioral pricing or some other thing, we So can all we need is a model that predicts markups test and intruments that are relative it.

to the change in markups that are excluded from supply, so for example, you can think of testing for cartels and you can think about vertical contracts, we are actually now in the paper and the revision testing marginal that's a model that specifies markups, we can test it, I will show you the example from the labor markups and a second, we can also think about behavioral, there's a lot of papers that have said maybe firms aren't doing the optimal thing, thing, firms aren't optimally pricing, that's a model, we can test. Let me give you some examples of world other papers doing those kinds of tests, so this Russell and scuderi and scuderia's mark.

I paid for free last year, they are thinking of labor supply and actually wage setting behaviour, they have data test, from an online platform that they will admit is not super representative, it's a very unique setting for tech workers, in silicon valley. So this is not a typical standard labor market, it's a very particular one, but it has a bunch of interesting features like, for example, a worker post an ad for themselves basically and has an asking salary and that ad will last for like 2 weeks, firms can place bids, they can interview them, and then firms can actually have a final salary, it's a lot of interesting data, there's this whole process to it, but what they end up doing is testing different models of wage setting behavior by firms, so they would be like you're offering marginal product labor, they test that against monopsonistic competition obscenistic competition in a compared basic about whether firms tailor job offers or not comments on the right here, here, what you basically see is they are trying to detect that correlation between two models, which model looks more misspecified while the red model looks more misspecified based on the test that they do which is based on our work Wilma MY1 caveat is this is a bin scatter of their results, every bin

scatter looks good, so, the raw data might not be so pretty, but this is what they got in So we actually find in their favor of wade's competition where firms setting.

offer wages independent of candidate characters. So they're not actually price discriminated to the candidate. And that's kind of their main finding, and that's they're presenting their main statistics from doing their tests are they strongly reject perfect competition, I actually think there might be a scaling problem because I don't know how you get to stat 54.84, however, I don't think perfect competition is the right model, so we will see. OK. So of applying this method that's in a different setting the next one I example Recent presentations she basically called up all her friends of other firms and said, let's raise prices, obviously very illegal, don't do that, in this case, there's a smoking gun.

So we kind of know the conduct, we know what's going on, very fun paper to read, the woman took maternity leave and cartilisation stopped for six months, it's actually really funny, so, again, this is kind of paper where you just plot the data and say there, you don't really need to do much analysis, but they do, what happens is there are uncardalized drugs versus uncartiled drugs. And here's the price series for the two, I wonder what happened right it's a real mystery so, what they do, though, is they posit a model they think the correct model of a standing is that all models law firms have competitive pricing and cartilage markets before the cartel formation while non members just best respond.

So they test this model if others will actually skip that here's our test statistic for testing their model versus competition non members complying member entrance does not comply, they reject all of those versus their positive model, so, and they are able to say, look, we know what we think is going on here, the non members are just best responding, I will just mention one more actually, I don't have a slide for it, but Sophie wing has a nice paper this room. I'm sure is very aware of what real page which is an algorithmic pricing thing for setting rent and she's using the test in that case to think about are they setting collusive prices, but I have 23 seconds left.

So let me just say in Syria. We see strong evidence of own profit rather than other pricing, give us some other model of common ownership if you want us to test it. We are happy to test it. And we are happy to see other people detecting it in exotic settings, my time is up, thank you very much. We have time for a couple of questions.

[>>]

[>>] Matt leiston FTC at this point. I think everyone knows who I am because I've asked 4 questions already, so this whole talk is about the pop style approach to demand 1st and then markups, oftentimes we have industries in which we can estimate demand, I'm I'm thinking about cases where we have prices and maybe we have marginal cost. But we don't have quantities, we have to take a different approach, intuitively, we are trading one variable Q for the other variable MC, but I wonder the prospects for being able to simply estimate conduct parameters or conduct models in settings where it's not so easy.

[Michael Sinkinson] That's a great question. I am bullish that's I would say we are starting to think about this as future work, we wanted to think about different games for this kind of testing, What variables could we apply this procedure to, so bullish yes, I don't have an answer for you exactly on the gasoline case, but I'm glad we are thinking about it.

[>>] Just to understand to make sure I understand, what you're rejecting is some model, you have to assume some truth.

[Michael Sinkinson] The null hypothesis is that the two models are close to the truth and when you reject you're saying one is definitely closer to the truth.

[>>] And then, the second question is can you then flip this test and instead, figure out the of markup mark that will be minimizing the correlation, That would be hard, in my head. That's akin to basically saying estimate those profit weights, so that's why Erwin nothing 'cause then maybe let's add one more parameter and see what we can,, so the idea of estimating kappa you have to impose a lot more restrictions than what we are doing, because, it would be a too many parameters problem, so a lot of papers that are doing the parameters they're estimating one, they're saying it is one it's constant overtime, bear with us, so people, this debate is actually still happening this should you try and estimate what that parameter is or should you try and test different models, so there are people who are going down that path as well.

Could

[>>] you explain why you don't use kappas for the test because it seems that that would be a natural variation that wouldn't affect marginal cost. But Sinkinson] they don't affect demand, that's why, the finance people actually the reason they told us not to use them is because they feel the ownership is endogenous to everything, which I think, I don't know if I believe that, but that was their point. So fewer controversies is better.

[>>] You take a very flexible approach to demand and supply through random forest, but PLP optimal instruments the results are basically the same, or the test statistics are basically the same, so I'm just wondering how you think about that, is the random force basically coming out as good as linear, is that specific to the serial setting?

[Michael Sinkinson] That's a great question. We are still trying to unpack it to be honest, it should always be better than the 1st 3 columns it has all of the information with the 1st 3 columns, our best guess so far is just the functional form of the demand optimal ID, the D side D theta, it happens to correlate well to the difference in markups that's our best guess. But we don't really have an answer. We are still trying to unpack it. That's the answer, it was a bit of a pleasant surprise. But I have to say now everyone went estimating PLP bothers to do that stuff, most people just throw in the BLP instruments and say they are done, it's just one more step after that.

Great, thanks so

[>>] much.

[Michael Sinkinson] Thank you Now we . are going to break until 10/25 for our last paper

[>>] session . It's a session organized by Mike sinkinson and we are going to start with Simon munghee who is going to present a much anticipated at least by me, paper on merger guidelines for labor market market. Mongey] OK, is that forward, OK, thanks for having us on the program, so, the agency's FTC and DOJ, the merger guidelines feature presumptions of anti competitor effects of mergers based on hhi, thresholds, the recently proposed 2023 merger guidelines have much stricter cut offs in that they kind of test anti competitive behavior they're and levels of hhis are smaller or changes in hhi's are smaller, it's tightened back to 1982, guidelines relative to those are in the 2010 guidelines, now these thresholds, and competition policies being developed and almost exclusively applied in the area of product markets, we don't really have frameworks and thinking about how we would apply these in labour markets for these useful thresholds for Labour markets, meanwhile, the 2023 horizontal merger guidelines and a lot of work recently in economics, which we have been involved in as well is sort of a lot about perfect competition in labor markets, in our previous work we used census like lehd, and LBD data to kind of assess the strategic responses of firms in local labor markets and found evidence of that.

So the contribution to this paper is to extend our existing work on competition in local labor markets to allow that framework against existing empirics that document the effects of mergers in local labor markets and negative effects on unemployment and wages, make sure our framework kind of matches up with those empirics, and then apply it to kind of study to such guidelines would be appropriate different guidelines, and our approach is to introduce multiplan ownership into a model of monopsony, extending our previous work, our 2022 paper, before we get to applying that for the question of whether the guidelines of the 2023 or the 2020 ten horizontal guidelines are appropriate, for labor markets. We're going to do a bunch of stuff, everyone I won't get time to go through a lot today, but in the paper we kind of theoretically characterize post merger labor market outcomes, this is going to be a model where jobs are differentiated and workers are choosing which firms to work at, it's going to be a model of differentiated, if you were to map it into a model of product markets you would be mapping it into a model of differentiated goods. It's differentiated jobs, look at labor market effects. We're going to make it flexible enough to have decreasing returns or increasing returns to scale and production, and their own existing results theoretically for the effects of mergers even in product markets where you kind of move away from constant returns.

So we are going to extend those but on the labor market side, we're going to show that when we replicate the quasi experimental conditions that we view as underlying David Arnold's really nice paper which is under review at the AER, in that paper, it kind of documents following a merger as an elasticity of employment labor markets, there's wages in labor markets they're higher when those labor markets are concentrated and we'll compare the model to his results and ensure that they line up very well when we hold it accountable to the conditions in the Pango random house for Simon Schuster, so the model generates predictions, which are kind of sensible, if you think the evidence presented in that case was sensible.

What we are going to do is we are going to assume a 5% gain to ultimate firms involved in a merger and then we're just going to apply the merger guidelines in the model, we're going to block mergers in market the 20/22/2010 guidelines, when we apply the stricter 1902 or 2023 guidelines, the efficiency gain offsets the labor market impact of the merger, and workers tend to be better off, when we apply the loser 2010 guidelines however, the labor market effects of the merger offset the 5% efficiency gains, and then we're going to kind of move to a way of assessing these mergers in terms of what we are going to call required efficiency gains when you come to me with labor market.

And we would tell you what we think they're required efficiency gain is that you should be asking your firm to kind of claim. So there's some concentration some changing concentration I'm going to need to see efficiency gain of 8% or 9%, these results are really sensitive to what we think for increasing and decreasing turns to sale, may be familiar with product function, which increasing scale and production when we replicate when you move over to an increasing returns to scale where mergers make work is better off under both of these situations, so having decreasing returns to scale in the framework is an innovation relative to what exists on the merger side and product markets.

When you think about his work in the ER having decreased scales is kind of important for understanding the impacts of mergers labor markets markets. OK, so just quickly two pieces where the labor comes up in the 2023 horizontal merger card lines, so, the guidelines are open to the idea that the level of concentration at which competition can concerns may rise, maybe lower in buyer markets than in seller markets, indeed, that's kind of going to be what we find, we're going to have a supply system or the

underlying parameters of that supply system would reflect deeper microeconomic behavior, and then across labor markets we will find very low elasticity substitutions across labor markets, a little lower than what you would estimate across product markets, on the product market side, that lower elasticity across labor markets means the mergers are going to have more harm in local labor markets, as per the guidelines, we are going to separate out, we are going to hold our nose when it comes to product market competition, and just focus on the labor market, so the guidelines say we're going to think about this separately, in the slides, I'm a little just assume that the product markets are perfectly competitive in the models we have extensions to monopolistic competition combo there may be a constructive way to think about this as you've got lots of local labor markets, the firms are competing legally for labour, but then they are selling these goods in national product markets, so these two issues are kind of separated. OK. Thomas. I'm going to give you an overview of the model, I will talk about this replication and the empirical evidence from Arnoldt paper, and then we will go through the application merger guidelines. And then if the I end come to the end, I'll come back to this penguin house for the Simon case, it's the key features of the models that has many local labor markets, across which workers are mobile, and workers are going to choose which labor markets work at, and which firm to work at in a particular labor market, and those firms are going to offer different late wages, workers preferences at a deeper level amenity differences across firms these are all going to end up encoded into elasticity substitution that we are going to estimate they're going to take those from our previous work.

And I will tell you how we do that in a bit, the firms are going to be strategic In our Arkansas paper we also kind of test this model, abstracting for mergers against other empirical observations, which shows that when the federal government says via hospitals have to increase the wages of nurses. They kind of estimate the spillover effects on wages of hospitals nearby, and then they find no effect on wages of hospitals further away. And we make sure that we test that it lines up when we replicate that empirical exercise when the model, that, again, would be very hard to think about if the firms were truly behaving atomistically, so I think when we take care.

Broadview of the data, using lehd LBD, census data, I think we kind of find evidence of this behavior in a lot of markets markets. Maybe not in some specific ones, bye. Think it kind of exists across the board.

OK. This is not a macro approach, and we've gotten this in the past that you guys are macro economists that you guys the macro approach, this is like up right right? So what we got to do is we got to talk to you guys about the microeconomic structure of what's going on, it's not maybe as rich as you would like, the idea is to nest all of this stuff across markets so that we can think about general equilibrium, we can think about welfare, because we are macro economists, so the labor supply system that we are going to have in this model is going to look like exactly what you have, in fact, we have a long proof in the appendix of Aer paper that shows you exactly that, and how to map the parameters of the models backwards and forwards to one another, so we want to be to able do micro LBDLEHD data to discipline the micro features of the model, but have a macro architecture.

So that we can about welfare welfare. OK, so I'll put more on the model it's going to have more labor markets and they're going to be indexed by J , there's an issue of how do you define a labor market, when we take this to the data we are going to do something that's kind of heroic, we are going to say the labor market is a three digit industry and a commuting zone, and that lines up with the way that David describes them in his empirical paper when he is thinking about mergers, this is a framework, came to me with the labor market Aiken investigate the model if you narrow down the labor market.

What you are going to find is that you're going to estimate higher elasticities of substitution across markets because the markets are more similar to one another. And because the markets are narrower you're going to find higher substitutions within markets, and at the same time, the same concentration of markets would go up. And what we find that's basically a wash, so as you expand the market the elasticities of substitution fall across markets and across markets, but concentration goes down, and you basically find the same kind of negative effects labor market power on the aggregate welfare.

But what for I show you today. We can do three meetings on, the markets and mobile across markets, they're not stuck in no market, and how mobile they are across markets is going to be kind of like key to our estimates, there's a finite number of firms in each market, we are abstracting from entry and exit, and what we do is we just take the distribution straight from the LBD data, in the LBD, what we do is we take all of the establishments of a firm within a commuting zone, and we aggregate them up, we compute total payroll, we divided by total employment to get a measure of the wage of the firm within the market, and so to fix ideas, there is 9% of markets in definition markets, 9% of markets have one firm, but like most of the the employment economy is in the markets with more than 200 firms, so most of the market firms most of the employment economy is still in a pretty competitive market, and there's only a small fraction about 1/4 percent that's actually in market where there's only one firm, so, markets still look pretty competitive, but we're still gonna monitor all of them in the data, are going to allow for increasing returns to scale, I think this is important when it comes to mergers because it allows you a real way to think about post merger, scale, in an estimate decreasing returns discount on the labor market.

But with increasing returns to scale made much lower efficiency gains to offset the neighbour neighbour markets and the mergers, the households are going to supply labor to all of the firms, I and all the and markets and that wage, problem that the household solves is going to deliver a labor supply curve that the firm is going to take as a given, that supply assistance is going to depend on the firm's wage, going to depend on the vectorable wages in the market, it's going to depend on other aggregate variables, the firms are going to compete. So it's like differentiated jobs, and the firms are competing in Cornell, the firm's labor supply elasticity's are going to end end up being smaller if the firm's small within the labor market, small.

If if the firm is large in the labor market and higher if the firm is smaller within labor market, it's going to generate some payroll share of the labor market dependence on markdowns, which is consistent consistent with Ev Profit maximization of the two firms jointly, so here. I don't have capital, which is supplied by households from their consumptions for savings conditions mama. So the firm is going to choose this competing core no, so the firm is going to choose employment, we do core no, you could do portrand, once you have differentiated goods. It doesn't really matter, you just get basically the same results out, am I saying something crazy no, when goods are differentiated.

They basically get the same thing, so there's one line of code which you can comment out and turn it over to Bertrand, referees don't seem to understand that. It's going to internalize diversion ratio when it increases employment of plan a, if they were to keep employment fixed by plan B, that plan B, by constricting labor supply in the market, a marginal costs are going to be higher, given marginal marginal costs are going to be higher. It's going to employ less labor common. So overall employment and wages are going to fall, the paper provides a series of propositions underneath Cordell, and returns to scale, which I think extend existing results from knock and whiston, groups for Cordova trend from

hydrogenating consulate, we can characterize the fair bit, not everything, we can show that the combined share of the merging firm falls and the shares of all other shares in the market increase come with this kind of crucially means when you're merging large firms, naive estimates of what's going to happen to the hhi kind of miss and can miss quite substantially, that's actually consistent with what David Arnold finds in the market in the data as well, a naive estimate of what's going to happen to the concentration in the market post merger is just add the up all firm the shares squared of free merger but because the merger firms are contracting their demand for labor.

And that's going to increase the market power of other firms and that kind of the increase in concentration is not as large as what you get from a naive estimate, which has implications for implementing the merger guidelines, we assure that all wages at all firms in the market of qualcomm you're giving more market power to the emerging firms, you're also giving more market power to the firms as well as the merging firms actually shrink, so, wages of all firms fall on the market falls as well. And you can show that at least one of the merging firms shrinks, this is testable of what's going on in the market following a He then merger.

required efficiency gains at the two emerging firms in order for average wages in the market to increase, we are going to try work towards a framework for thinking about merger review, which we are going to state in terms of required efficiency gains rather than these hhi and delta hhi cut offs. Again while we thinking about were these cut offs it's not because we think they are optional, but that's because that's the frameworks that's being used to analyze mergers. OK, we talked a little bit about how we calibrate and how we validate the model against this paper, so we're going to take the calibration of the model from our previous work, what did we do in So as that I paper?

supplied that if you're a large firm within the market your labor supply curve is less elastic, it would be great to be able to estimate those labor supply elasticities, plot them against shares against the labor market, across markets, the reason is because the markets concentrated, going to be responding to one another, so if I respond to markets firm, teammate, Indirect inference, so that's what we do in that paper, what we do, is we take a firm, which has plants in Minneapolis and plants in Duluth, which is a place in Minnesota, I didn't know where it was, so I've got plants in both of these places, I'm a firm, I'm small and run market, but I'm bigger than the other, then what we're going to have is a shock to marginal revenue going to hit both of those, across market variation and how employment wages respond, to different shocks to marginal revenue, was kind of more elastic.

And to lose more at versus Minneapolis where it's large. so that we can't use these and go directly to the straight structural parameters that we like, that we want, we can still use this as an exercise, which can kind of inform us as to what these are and practice what we use as changes in state corporate taxes, which are idiosyncratic in the sense that C crops and S crops, so it's not a pure aggregate shop from which we would learn nothing. So that's what kind we do in that paper, of these deep parameters which control the parameters across the markets, the substitution across markets is very low, which is consistent.

If you think about papers like in walpen, and always back out in models of labor mobility across regions at this very high cost of moving across different places, and that's kind of what's encoded into these low elasticity substitution, are those high elasticity substitutions within markets, and that's kind of what you really need to rationalize that employment looks like it's very mobile across firms within a market when relative wages change, so that's kind of key. I think in different product markets, we tend to estimate that

those parameters are kind of closer together, we also estimate there's decreasing returns to scale,, scale, which again makes sense, which when you think about doing this at the lower labour market level as opposed to the national level where you may have headquarters and things which lead to increasing returns of scale, OK, so that's kind of like we estimate the model we then show that the model replicates this evidence from David Arnold's market paper, so what does David, do.

He assesses national mergers, he takes a national merger and he looks at how employment and wages responded different labor markets, has segmented by the concentration of labor market so they merge they're present in many labor markets, they're big in one small and another, once concentrated one's less concentrated, you think that the national mergers is kind of orthogonal to the characteristics of any of these particular labour markets, and he's going to compare how responses vary, now, how do we stimulate this in the model, we're just going to simulate like 200,000 markets, on average, firms that merge.

Are large, so what we are going to do is we are going to draw 2 firms at random, from every market, and we're going to throw them back we are going to keep them and merge them, if they are pre merge or employment is above some cut off, and what we are going to choose, we're going to do is choose that cut off. So that average the average size firm of emerging firms is roughly the same, it's exactly the same as what he has in his sample of merging firms, I will tell you how we would like to improve this as we go on, so as we the firm's initial size is we keep them.

And that's going to deliver exactly the same medium size of merging firms that David has in his exercise, right, and we just merged the firms as I showed you before, we do the equilibrium of the market, and then we run like apples apples to exactly the same regression that David is running in his job market paper, so we have the same medium for market employment pre merger of 116, which is big, right, if you think the average firm size in the US is like average firm size is like 20, and I'm talking. This is now like within a market, so these are big firms that are merging, so it's important to kind of subset sub sample the data in this way, so what David found was that there was large changes in employment, so 14% decline in employment, and a .8 decline in worker earnings, and the model is pretty consistent with that, so again, this is because I've got this high elasticity as a substitution within the market, we can get large movements in unemployment with changes in wages wages Nonetheless it's going to be small, which gives substantial labor market power, if you .

were to try to rationalize this data with low elasticity as the substitute like a monopoly model, where everybody had the same elasticity of substitution that was the same within an across markets, then you wouldn't be able to get these large employment changes at the market level, or if you did, you would infer that there was no market power because it would just look like they were very elastic, OK, also, consistent with what David finds we find that these effects are effects on worker earnings, they're higher from markets, there's a lot change of 3.1 in worker earnings, and finds in high concentration markets, which he defines his high concentration in terms of exactly the way that the horizontal merger guidelines define them, and we find that that's kind of a factor of four higher than low concentration markets in the model, why is it higher in the high concentration markets, if you merge 2 firms in a very competitive market, and market share doesn't really increase pretty much, the market share of the other firms doesn't increase very much,, there's not much of an effect effect on wages or employment whereas if you're in a concentrated market, that's not the case, so, a model generates is a data generating process, which is consistent with.

I think the best evidence that we have on the local labor market effects of mergers, and so we are going to say OK, I think that's alright to kind of proceed with,, to assess the effects of different merger guidelines. OK, so what are we going to do. We're going to simulate the emergence under the material, under the simulation of David Arnold's paper, so I'm only going to keep firms above a certain cut off. So we have kind of the same empirical properties of emerging firms at least in terms of size, that we observe in the data, and I'm going to assume what our co author says is like maybe a standard assumption of an efficiency gain of like 5%, I'm going to increase the productivity of both of the firms in the merger by 5%, and I will show you what happens if I choose different numbers, I'm going to pass the merger in each market.

If it satisfied some guideline assessed at the market level, the way that I understand my co authors understanding of how the law is applied, is that you can you can basically treat this market like market by market, not looking across all these different markets comm that as a comparison to an airline case, you'd be thinking about this route by route, he would be thinking about it like labour market by Labour market, and so we're just going to apply it like market by market, and then we are going to compute the welfare gain slash loss in each market, and then we are going to average across the markets in which there is a merger, we are going to do this in mid the 2023 guidelines, which are tighter, where we prevent mergers above a level of an hhi of 1800 and a change of 100, relative to the 2010 guidelines, which are substantially figure, and what we find is that these are dollar values per market where there's a merger, we find that there's welfare gains in markets with permitted mergers in the 2023 guidelines, but on average, underneath the 2010 guidelines there's welfare losses in the markets where merges are permitted, that's kind of the guts of our conclusion, that workers are harmed underneath the 2010 guidelines, whereas underneath the 2023 guidelines, the efficiency gains associated with the merger are translated over to welfare improvements, of the worker, if I shift this to increasing returns, so again, we have decreasing returns, then here's the key typo there's not supposed to be a minor sign in front of that 112,000, you know, this is where typos are the only way for you to check your economic intuition and your code.

That's supposed to be a positive sign as well, so if I will go over to increasing returns and scale. I would find welfare benefits in both cases, why, because you get the gain beyond an removal of the employment to decrease output, and that offsets some of these labor market consequences, so in the last minute, you know, maybe these a different way to kind of think about this is through. We're trying to think of it as required efficiency gains, so if I define by delta as far, like a merger by merger level, a percent increase in firm productivity is such that worker surplus is mutual, which in our case, is the wage index and the market level remains a constant, and then we can kind of stand on the permitted urges and literature the 2023 guidelines, so, the efficiency gains are just below that 5% cut off, so if you think what the actual efficiency gains of a merger are 3%, and even under the 2023 guidelines, that wouldn't meet the standard of how much efficiency, you need to offset the negative labor market effect come under increasing returns to scale, this was much lower, or again, this might be the appropriate standard in product markets at the national level or we think that there might be increasing returns of production, but locally in labor markets, that's kind of not the case.

OK. Comm these require, they're like a distribution of these required efficiency gains underneath any kind of observables underneath any particular merger, and so the guidelines are in terms of changes in the concentration and levels of concentration, and so, I can kind of take all of the mergers that satisfy the cell

in terms of changing concentration and the level of concentration, for example, if I take this one or I have very high levels, large entries and concentration come on a high level of overall concentration, and the median required efficiency gain there is like 27%, which is like very large, so you're you can kind of stare at this heat map and here go OK, here's what I think if I required efficiency, here's what I see think efficiency gains kind are, I think they are around 7%, if you I think they are around 7%, then you go OK, of half of what the time I'd be right, letting through mergers which 200,, have a changing assumed concentration of about 100 concentration about efficiency gain is about two, 3% what you think is kind of like the the standard, and if you impose a rule of assessing all mergers change in concentration is around 50 to 100, that basically gets you all of the way there, you panopsy are kind of capturing all the mergers that you think would not be amount, deliver welfare improvements ID efficiency gains like not offsetting the negative labor market impacts, OK, I will skip to paying a random house murder because I'm totally out of time, so, I think the paper provides quantify analysis and merger guidelines in labor markets are taking a model of all adopting labor markets and extending it to look at mergers and then testing it against empirical evidence, on mergers, and the key result is underneath what we would assume is standard efficiency gains.

Workers are beyond harm underneath the guidelines, which have substantially cost tightened relative to the 2010 guidelines, and the paper we provide down with wage formulas, and how these relate to required efficiency gains in the distribution of gains and losses across these different observable characteristics, thanks to Matt

[>>] OK great, and now to discuss we have Matt Weinberg from . the Ohio State university All right, thanks to the organization for including me and for the sponsors and Weinberg] assignment and his co authors for writing such an interesting paper it was a across pleasure to read it, how does this trying Got work. OK, so given all it. the recent interest in the potential labor market consequences and mergers and the fact that between 1:00 and 2000 are proposed in each year under the ASR act, it's really useful to think about screens under help splitting those out that would be problematic for those who would not come which is just what this paper does, so it's an important paper, it's building off of this previous paper by seminars co authors that tries to do something that I think is really important and interesting and ambitious, it's the entire US economy, so, it's going to require a finding every labor market in the US and thinking about of labor supply system that would govern how workers would choose which job to work at given the options that they have, and the model of competition.

So it's a tough problem, this paper builds on that, by allowing for multi plant firms, so you can study mergers, so the framework here is really great, it's allowing you to do kind of large scale merger simulations, you can kind of simulate any merger that you would think about in the economy, and that's used to evaluate different screens, you can use it to evaluate any screen that you can think of, because you can simulate any merger that you might think about, you think about whether that would be caught by that particular screen or not, so the paper shows that if you assume fibrosum prodigains against mergers, and apply the 2010 concentration metrics, to find for product markets, on concentration measures competed off of wage bills, then, the average merger would reduce wages, but it wouldn't under the AB2,, so what I want to do in this discussion is 2 things, first, I want to just lay out how this framework relates to something that a lot of people in this room.

I think would be more familiar with kind of the canonical merger simulation, and 2nd, this is kind of a tough criticism I think there's a lot of value in doing this with scale, and using this framework to look at

specific cases, so, when I think of the canonical merger simulation framework, I think of, for example, the 2000 Ram paper that you guys probably all know most people probably have looked at this very carefully, it's just a few key ingredients, the 1st is it's assumed that there's price competition amongst firms selling differentiated products, and, there's multi product ownership, so when firms own substitute products they're going to charge bigger markups because they can internalize lost substitution rate, and, we've got this.

Nice work of writing out the markets that bring out this one model, the markups are going to depend on ownership, it's characterized by this ownership matrix Ω , demand conditions, and apologies for the typo, of course, there should be an inverse around that matrix, it's pre multiplying the quantity of vector on the right hand side there, some of the Protran game, of course, you could do this for core no competition too, would that change this, well then, there wouldn't be a typo, there wouldn't be an inverse there, but, you would be taking the Jacobian at the spectrum of the inverse labor supply matrix, so, Crystal Levin and Waldogel is not this Waldogel but the other one is some co authors have a really nice paper showing that that you the markups would get under Cournot with the exact same demand system and cost are actually bigger than what you get in for Trad, and they have implications of how it would work too, it's an interesting implication, so, this paper has a map into it, so I'm going to do a simplified version of the model, actually more general than this, they did a good job of discussing, but, so now, instead of output, price competition, we've got a perfect competition on the output market, the input market there's going to be imperfect substitution of workers are going to view jobs as differentiated, and, if you assume constant returns to scale, you get this simple profit function expression right, the firms are going to choose how many workers to hire at each plant that they have, to maximize their profits, the value of a marginal product of the worker, and constant returns to scale, minus the cost of hiring the workers scaled up by how many workers they hire, and you can rewrite that same system of first order conditions that you get out of the Cournot model to look like something kind of familiar with a similar expression to what you would get of what people have looked at already, the markups would just depend upon the labor supply less derivatives and ownership, OK, so, what's the merger do, what's going to change ownership, it's going to potentially change productivity, so, I want to dig into that a little bit now as well, so let's take a simple case, where you've got two single product firms before the merger, pre merger each firm is going to charge a markup that will depend upon how inelastic labor supply is to its firm, and, after the merger, you get this downwards wage pressure term that accounts for the fact that if they want to hire more plant, it's going to increase the wage you have to offer a plant cake, hello, the key question is what does that look like right, so this paper has again, a labor supply system that you can do at scale, so it's going to be pretty simple, it's going to depend upon two key parameters and the last substitution within markets and one across markets θ , so given that labor supply system, you can write down without downwards wage pressure term would look like, that's going to depend upon those two parameters in the wage bill share and the wage listed, OK, I just want to offer two comments, the 1st is that in the spirit of Mike's keynote, I think it would be interesting to think more about other models of wage setting, an easy one was Simon mentioned and I've kind of mentioned already would be the wage study just posting, that would be straightforward to do, and, I think the papers got a really nice framework for trying to think about that these two different ways to model it, which one is best, you could think about which framework generates data the best replicates what David Arnold found in his other paper, and then the next question would be, does it matter for what happens with mergers?

So straightforward, that you could come up with other stuff, maybe more realistic it doesn't apply wage setting within a firm, stuff like that, how you might implement this on specific cases, so note that downward wage on plant one after a merger with the firm that owns product 2, it just depends upon product firms' shares. There's a lot of data out there. Alex Masters' work, for example, that shows that workers really have strong preferences over different menus: they care about hours, they care about commute times, they care about autonomy in the workplace.

And so on, almost like the DD assumption but that's hard, when you start thinking about jobs as being differentiated, I think about us all being differentiated from this BLP framework where you think suppose I knew the different options that are out there, attributes of the different options, maybe attributes of the workers, and where people go, the tough thing there is if you've got like a lot of wage dispersion across different employers, I think that it's often not the case that a worker can get any case job that a worker can have at one time, so it's all kind of logic shock, you're gonna get super inelastic supply.

But it's not because people don't care about salaries, because they can't get any job they might want, so hard, another thing that you might do though if you've flipped this around and think about the patram model, that downward wage term is going to be like this to anybody in this room, going to depend upon the markdown and divergent markup ratio, and maybe that's something that you might be able to get a little bit easier under time pressure, so, what's the diversion ratio going to become it's going to be the fraction of people that switch to plant one on plant 2 reduces its wage, so, oftentimes when workers leave firms there are exit surveys and they ask them where they are going, stuff like this, so if you have that kind of information, you can figure at least something, it's not exactly how people are switching given all the equal wage changes, you might be able to use to inform this diversion ratios which I think would be interesting, so I'm curious about other ideas the authors might have about that kind of thing, I know that's moving away from being able to think about aggregate issues, which I think are really important.

But these are the kinds of things that I think will come up in specific cases with this sort of framework. So that said, I really enjoyed the paper. I encourage you guys all to take a lurk thank you for the opportunity to be here. Thanks, Matt for the discussion, we've got a paper which we use like Norwegian

[>>] I don't know if I talked about this on the talk

[Simon data. Mongey] And this idea where we use how workers are flowing use general are flowing Norwegian across firms or professors are flowing across firms. And you can define labor markets more broadly like our labor market might be natural across many different locations, whereas other workers labor markets might be local, the US is that we don't really have that data publicly available to us, but as Matt kind of points out, I think when you think about what you guys can access here, when you can subpoena data from firms all those exit interviews. I think that's exactly what would be incredibly useful, for occupational data at the firm level the best we really have in the US is the sense this, but you can't do anything about flows in that because that's just there is better data in other countries, but I think that's exactly what you want to cross be doing in terms of subpoenaing information, section.

Hi, I was thinking some of the merger propositions could be improving efficiency of

[>>] management in reducing payroll in getting redundant worker or overpaid workers, how do you go about thinking about this entangling these in the effect from restructuring versus labor market power,

[Simon question. Mongey] I think looking at that's things under decreasing scale a gives you some sense of that because great increasing returns to scale we think of it as literally recording that dependency, and merge the two plants and get rid of a bunch of workers and still keep the same level of output, and can

we show that that's that's important, a reduced form structural way of thinking about it, rather than, yeah, again, if I were to think about David Arnold's paper, it kind of again runs up the opposition we don't have that occupational data, so we can't see he's using the LVD and the lehd, can't see what the occupations of these workers are, and who's getting laid off, either that form firm or in the rest of the market where you also find negative effect, so, the data constraints are kind of real and thinking about that.

But it's a really good question, I'm wondering how you should think about the endogenous mergers I . think all of the mergers are random,

[Simon Mongey]

[>>] slide where we're trained to approve Yes, this is incenses we getting have better the data, variance covariance structure of employment within markets amongst emerging firms comment could be the case for all the firms that are merging relates to east coast air conditioning or west coast air conditioning and there's no overlap between any of the firms at all, this is like going halfway to answering your question, so we are getting better data on whether you know the firms that are merging within markets are both big big and or one one's small etcetera, so we can kind of use that in this additional screen on our merger simulations, the next step would be then having a theory of why then the firms would be merging and that would be endogenous itself, I think in terms, this is really nice about David's paper, the merger itself was endogenous, but at a national level and you'd really worry about that from measuring the product market effects of a merger, but then I think the fact that you can use the idea that the kind of concentration, the heterogeneity confirms within each local labor market somewhat orthogonal, I think to say why the firms are merging nationally is like a useful insight, you can take care.

National merger, compare the firms in markets where they are both big, both small, and learn something about those dodging concerns., of course, in the model all the firms want to merge because it's going to increase profits and all of them come out. So like leading papers where we have those. And they are endogenous, it usually ends up looking like a random encounter where you then go OK, I'm going to merge which isn't like much of A richer theory of mergers, so I think that's something broadly to think about, in these models we always want to do it, Have you thought about the role of managers where all of a sudden, they could be managing different plans after the

[>>] merger and these kind of affecting the way in which kind of high skill labour is organized and do we want to think about this as we study these Yes but again, the limitations we questions? don't really know a lot about it, kind the of the same answer as to Leon, having some increasing returns to scale. I Mongey] think is a way of approximating that in the model, you can make some workers redundant, and still keep the level of output the same as it would be by reallocating production to more production plans, increasing returns to scale even if you didn't have any heterogeneity and jobs you'd still want to allocate all of the plans because you've got infinite product as you drive each of the plans to 0.

But the increasing returns to scale means you really want to shuffle everything to one plan, and underneath that, you think I'm economizing on management overhead or something like that, but you still in this case want to operate keep on operating both the plans even with increasing returns because you kind of drive the national, and you can pay him very little to work there so, that's how we think about that in reduced form way, yeah, I think it would be great to be able to dig more into that, just kind of get into more data on specific mergers, and how they end up reshuffling management Thanks for the last question As mergers become more profitable and we can see

[>>] how that can be passed through to workers, and we try to rationalize those moments with the model

[>>] I guess empirically it seems like they don't, in that again, I'm leaning on the Arnold paper, wages seem to [Mongey] fall now, and this kind of relates to catcher's question if that allows you to expand your scale as a manager, maybe some of the workers in the firm would be getting compensation if they are kind of in competitive markets, just getting this back to the other question of maybe this is different for workers and different occupations within the firm, and certainly Hottel Williams work shows that productivity shocks the firms in terms of innovation have different consequences for different workers within the firm, whereas here we are kind of looking in the wage, the application of the merger guidelines, especially as written, and enforced, but, yeah, you want to kind of split them out by occupation, and look at the heterogeneity and those effects of those workers, Great thanks.

[>>] OK, next we have Brett Hollenbach Anderson from UCLA School of Management Anderson of management All right, thanks a lot. Thanks to the organizers, so this is joint work with my colleague Ashwin Fountaine, he wishes [Hollenbeck] could be here, he's expecting the birth of his first child literally any second now, but he sends his regard, so, this is about fake reviews, we take as our starting point, fairly well documented in fact, by now which is that sellers reputation in online I have a large causal effect on their outcomes and their bosses, and prophets, so particularly sellers have a pretty strong incentive to manipulate the methods and as we saw yesterday, platforms are not liable for their sellers actions in most cases, and, platforms want lots of sellers to be active, so it's plausible that they are under regulating this type of thing, as a consequence, Brady manipulation is extremely widespread in E commerce comments arguably a bigger problem now than it ever has been, we believe this last claim is supported by our research, we also believe it's put out by the FTC recently, and, we have reasons to think this is bad, potentially, we know that online ratings and reviews are a key thing that make online marketplaces work by solving the asymmetric problem between buyers and Sellers, and we have good empirical evidence from the same setting that we are studying, which is Amazon reviews, that these have large benefits of consumers, so probably not surprisingly consumer protection regulators are seeing this as a large and growing problem almost of the FTC our host has put out rule making just this summer trying to define and codify different types of radiation.

Radiation practices UK competitions and markets authority have done something similar proposed law and doing the same thing, and it's also done something similar as well, so the goal of this paper, then is to study the impact of fake reviews on Amazon on Amazon sellers on the platform itself, and then you know most notably on consumers, so what we do is we provide theoretical framework for the ways in which fake reviews can impact consumers, and then we gather there are two types of data the 1st is a data set of what Amazon products are using fake reviews and when, in the second is a set of incentivized survey experiments designed to elicit consumer beliefs about how common fake And then we take those data, and use reviews it to estimate a structural empirical model that are.

captures the main forces at play, and then use it to simulate counterfactuals where we remove fake reviews from the marketplace and compare outcomes for sellers and consumers, all right, I'll just dive into the model, so we think there's two main ways in which rating manipulation or fake reviews impacts consumers, at first just what we would call misinformation, that is, fake reviews mislead people. And cause them to purchase products that they wouldn't otherwise purchase, and the main way this manifests itself is people are buying lower quality products than they would prefer to, a secondary effect of misinformation is that the products that are using fake reviews can also increase their prices because they have higher ratings, so even the people who may have been happy to buy the low low quality

product might be harmed if they end up paying a higher price because of it and then, countervailing effect might be that honest sellers those who are not using fake reviews are in more intense competition with those products come and therefore they actually have to lower their prices, which could benefit consumers, and then, the second channel is what we would promise.

Trust which is that in the long run, as large numbers of sellers use fake reviews, consumers lose trust in ratings, these ratings become less useful for solving the asymmetric problem, and demand is less sensitive to ratings, this again, could manifest as people making purchases or choosing products that they wouldn't have otherwise, but again, there's a countervailing force here which is that as ratings are less precise signals signals of quality, and price competition increases as a result, and people can be better off because of that, this is a simple model. I'm just going to demonstrate the main mechanisms visibly 101 demand curve, consumers get some surplus, producers get some profits, not reinventing the wheel there, but, what if a firm uses fake reviews, they should shift their demand outwards, if demand is increasing in ratings, which it is, and, consumers are going to buy more of those products, they're going to expect to get more consumer surplus, than they would have and the consumer surplus that they actually get is the realized quality, not the quality that they expected, and it includes this disutility this red triangle that is reflecting the fact that they bought too much of this low Alright this is profitable for Sellers, and both of quality these results are just holding prices fixed to come product.

of what we would expect that if your demand shifts outward, then you should increase your price, and starting from the previous consumer surplus, we moved to this one, so in a simple linear demand case, the red triangle dis utility is the same size, but there's an additional welfare harm in the size of the gray area, which is reflecting the fact that prices have gone up for these products. So even people who would have bought them otherwise. Now are harmed because they are paying higher prices. All turning to the honest seller or the sellers who right are not using fake reviews, they are now in more intense fella competition, so there is a dual demand curve that shifts inwards, and so, consumers were purchased at that level and receive consumer surplus that is actually greater than what they expected, so they will actually end up greater quality product than expected, then the size of a gray triangle which is that they are not buying as much of this high quality product as they would have in the absence of fake reviews, but again, this is holding prices fixed, in reality, we would expect when your shifts inward.

You lower the price, and so the consumer surplus for these sellers increase actually as they lower the price, and under reasonable conditions people are going to be better off to buy these products than they would have in the absence of fake reviews, so the people buying the high quality honest products could actually be better That's misinformation when we look at mistress, the main effect off. that we think is likely to occur is that products with high ratings, are going to have their demand shifted inward because people don't trust their high ratings, they think there's some probability that those are caused by fake reviews, so they're going to put less weight on it, and buy less of those products, could also be the case with the low quality products that compete with products that have high ratings see their demand shift outward as well, because people are less sensitive to ratings, and in general while a lot can happen here, relationship between ratings and demand breaks down a little bit due to this mistrust, and we think the main thing that's going to occur is that again, people are going to buy a set of products that they wouldn't have otherwise.

But it could be the case that price competition is increased as a result, so turning to the data and the setting, the setting we look at is Amazon, and the main channel by which Amazon sellers get fake reviews

are these large private Facebook groups and the way these groups work is that the Amazon seller or some intermediary working for them post a picture of the product in a brief description and an interested reviewer will reply to that and say, here's proof I have a valid Amazon account and everything common and that reviewer actually goes and buys the product, so these are verified purchase reviews, they post the positive review it has to get past the Amazon's initial filter, and then the seller will reimburse them the cost of that product and in some cases, a cash commission.

So this is what these posts look like say something like need reviews, you have to be in the US, I will Paypal refund you and cover Paypal fees send me a message if A note that is important to you distinguish between fake and incentivized reviews, are interested. 'cause there was some payment or discount in exchange for the review, and the incentivizer has the ability to leave a negative review and they won't be punished by the seller for that, actually, we have good empirical evidence for taobao a similar type of market that incentivized reviews can increase the efficiency of a market by causing higher quality reviewers come.

And then they get higher quality reviews in our case fake reviews come there's no disclosure, there's a strong requirement that the review be positive in order for the reimbursement to occur, so this is an important distinction. We observe these Facebook groups at a high level for about four months, there's a lot of them coming there very large there's thousands of members and then, there very active there's hundreds of clothes per day, we do a crude back in the envelope calculation trying to account for the amount of overlap and post between these groups, it looks like about 4 million distinct products posting in these over the course of the year, so a substantial number of Amazon products are using these So it's nearly impossible to scrape Facebook, so instead groups.

of trying to, we hired undergraduate research assistants to go into these research groups and hand collect data in a fairly labour intensive way of what products are in those groups buying fake reviews, and the result was an example about 1,400 unique products and we observed both the start and end date sort of roughly speaking of their fake review campaigns. And that with a very large then, scale daily scraping of Amazon to we capture product outcomes, pair the main way the scraping works is that we look at the keywords of the products, and we scrape everything that results in the keyword searches, we see all of the products, prices, ratings, ranks and so on, and then, for the fake review products.

The focal products as well as a set of close competitors. We set extremely detailed data on their reviews, and on the reviewers of those reviews, the people writing them, we have sales ranks for all products and then for a subset of products, we observe sales in terms of quantities. All are these products that right, are using so We can look up what categories fake there and they're not highly represented who or highly clear categories they seem to be reviews? coming from all over in a fairly dispersed way, if you look at their characteristics, again, it's not necessarily these are products with very low ratings or very few reviews, who are just trying to get started, they are a bit younger on average, and have fewer reviews on average than the other products, but the variation in each of these measures is quite large, and so, what we should think of is while they are a bit younger on average,,,,, we see a very wide swath of Amazon products that are All right, I will just note engaged I have a previous paper with other co authors that in presents a very large number this.

of mostly descriptive results about these products and how this market works and what happens to fake reviews, so if you want more detail you can look at OK, but turning to the structural model. What are the pieces that that. we need? So for this information, we need a model that maps people's observed reviews

to their beliefs about product quality, and a model that maps people's expectations of product quality And then for mistrust. What we need is a model of how these beliefs of into product quality vary with respect to beliefs about how common demand. our fake reviews.

And then like any common structural model we know we need to have product structurality, I'm gonna take a fairly standard approach to that the goal of it is to characterize how consumer inform expectations about product quality from an observed reviews, we need Three pieces we need a sort of model of how ratings are determined as a function of product quality and the use of fake reviews, we are just going to use model assumptions for this, we need beliefs about how prevalent make reviews are and what types of products use them are we're going to end up using surveys because we could oppose those assumption for.

This is what we could assume that people know roughly what proportion of products are using fake reviews. But not exactly which products, I don't necessarily believe that people have rational expectations and it's at least interesting to go out there and measure what people actually believe that's what we do instead, and then we need a prior about the distribution of quality, which we are going to estimate directly off So, the model of beliefs is of that products must improve underlining quality Q between zero the and one, and the way to think about data. this quality is that the probability that a product receives a positive review is higher this quality products are more likely to get positive reviews.

value And Q . products they are using fake reviews and they get So then an additional set of positive reviews from that, so for data θ is the portion of some their reviews that are faking that proportion is positive by definition, so we get sort of a nice model of what is the likelihood of getting a positive review as a function of quality of take reviews, and then we can say, well the distribution of ratings, the number of positive and negative reviews is going to be distributed by nominal meal that's coming out of this probability, and what this buys us is probability distribution over product ratings as a function of their underlying quality and whether or not they are using fake reviews, so we can take that, and then incorporate it into simple bayesian model, which we are going to use in the posterior that consumers use to evaluate expectations about product quality, so the first component I just described and then the other two components that I need are beliefs about the probability that a given product is using fake reviews as a function of its ratings distribution and, the underlying prior distribution of quality.

So that last one, we are just going to estimate based off of organic using the data, this is what it looks like, so this is the prior, and our estimate of it is addressed that products are on average lower quality than what we call non fake review products products that or are not doing this, and then from beliefs, what we need, is a measure of the fraction of products that people think are using fake reviews come up both overall and a function function of their ratings distributions, we also like to know the fraction of the products reviews that are fake if it is using fake reviews, so we're going to set up incentivize survey experience to elicit these beliefs, we run this survey on the platform prolific, we start with a sample of about 750 respondents, and, we get a final sample of about 400 after eliminating people who fail or various attention and comprehension checks right, the primary task on the survey is that people first select a set of categories that they shop on Amazon and, and ask them basically, well what do you think is the probability that this product is using fake reviews, so this is what that survey task looks like, so the product page that we showed them is the same one that appears on Amazon, and it's actually the way the product page appeared at the time that they were using fake reviews, we have the underlying HTML code.

This is from our initial scraping of this data, and so, we can modify the HTML code and change the average rating and the number of reviews that appears next to these products, so that gives us nice clean variation in things like number of reviews and average rating, that is sort of separate from the individual product fixed effects. So then, we ask them how likely do you think that this product is buying fake reviews, they move a little slider as they move the slider they get these clear statements of how their payouts vary. So they know that we know the truth about each product, and then they know that they are going to get paid based on how accurate they are, and so, they have a strong incentive to give their best possible guesses to the product that is using fake reviews so, these are the results from that aggregated cross product and respondents, to results come out of this, you look at this distribution for products that actually are using fake reviews and those that are not, they are almost identical, some people are very bad at guessing which products use fake reviews, and the other is that their average response is about 40%, which is surprisingly close to our estimate from our data which is like 35%, so rational expectations would work reasonably well here, so, we also look at how these beliefs vary with the ratings distribution, so people are more suspicious of highly rated products, there's not a strong relationship with number of reviews, we can look at the whole distribution, this is actually what we are going to incorporate into the model, we see again, that people are specially suspicious of products with a few ratings and positive ratings, we do a similar task for the proportion of reviews that are fake using fake reviews, here people are a little bit off, they are a little bit naive and that their average guess is our estimate of this value is slightly over So karma we then model demand, this is mostly pretty standard, 50%.

we have product fixed effects weak fixed effects limited set of product characteristics, and then instead of putting ratings or reviews directly into demand, what we do is we put the quality and demand that results from our model at the leaves, so we incorporate some features from the survey, and incorporate our own estimated fixed effects. And we use a model to compute the expected quality for each product given its ratings, all notes on demand commonly allow for heterogeneity and weight right. placed on quality, so a random coefficient on quality term, which Other is normal, we have nesting terms for each submarket, no direct interpretation of demand estimate and not as precisely estimated as one would like, but that caveat this is maybe a work in progress, we can look at the elasticity's, the elasticity's for price are a little on the inelastic side, maybe somewhat surprising, although consistent with other people who have estimated demand on Amazon, and then the quality of elasticity is strong, so common.

We can use that model to evaluate counterfactuals, and then the goal of these are to simulate in the market with and without reviews and use that to evaluate the difference to sellers and consumers, so we are going to start by removing fake reviews and how ratings change and how beliefs about product quality change, and then we are going to get the full equilibrium outcome in those two scenarios coming then we are going to decompose the channels of mistrust and misinformation, so, we have an estimate of the proportion of each products reviews that are fake, and then for the products using fake reviews, given the fact that those are all five star reviews just remove that that five star review that we see that those product ratings shift substantially to the left, and we look at all the products we can look at how beliefs about those products quality change in these scenarios, so surprisingly for products that use fake reviews from where they have substantially higher expected quality when those fake reviews are present, and then, for products that are honest or not using fake reviews, their expected quality is actually lower with fake reviews present, and this is sort of resulting mistrust,, consumers are sort of mentally discounting their

high ratings, and so, their expected quality is lower when other products are using fake reviews, then, we can look at how outcomes change.

So we take as the baseline world world is the we look at how things change when those fake reviews are present, we see that sales increase substantially for the products using fake reviews, they go up by about 30%, sales fall somewhat for the non fake review products, and the overall effect is an increase in total sales, I see pretty large changes in prices, so the fake review products increase the prices on average about \$0.70, the non fake review products cut prices by an average of about \$0.35, interestingly, the overall effect on average prices is basically neutral almost perfectly cancel out although the distribution of prices changes quite We can look at profits as well so this looks profitable a perhaps not surprisingly profits bit.

go up by about 35%, and then finally, we look at welfare and we find a net negative effect on welfare, so they are on net harmed by this, purchase price will note that there's an asterisk on this slide, this is preliminary, this is really a note to the regulators in the room, the magnitude of these welfare effects is potentially subject to change. So don't put it in any government reports just yet, while the magnitude of his effect is something that has changed, I think there's some interesting qualitative conclusions that we can draw from decomposing the channels that are not really subject to change, so, first, we take again, as our baseline in the world without bank reviews and then we evaluate two types of counterfactuals one in which we put the fake reviews back in the market, but people are fully trusting and and then ratings, the other in which we do not introduce the fake reviews, there's no fake reviews, but people think that there are, so that isolates the mistrust effect, we see that mistrust information and consumers harms consumers that welfare, lowers that effect is actually much larger than the net effect, and the reason is mistrust by itself itself is actually beneficial to consumers, so consumers benefit from being mistrustful of ratings, and this is kind of introducing an analogy information disclosure, which suggests that platforms might be better off not fully disclosing information about product quality if people have sort of less precise product quality, this increases product competition in a way that products can be positive for consumers, so we find something similar here, when we look at the platform's profits, we see basically the opposite effects, which are that the platform definitely profits from this information, so they profit from the presence of for fake reviews, but they suffer from mistrust, some mistrust lowers Amazon profits, and so we can think about this as a short run benefit to Amazon, allowing fake reviews is that this misinformation effect comes in their profits go up, in the long run, there's a cost for that, which means people could become less trusting of ratings, and made by fewer products of Amazon and might lower total revenues and total commissions that Amazon receives and if we sort of look at the full row, we see that Amazon would as far as to get rid of fake reviews, but if they just deleted the fake reviews, they would lose profits because people would still have this mistrust, so they would have to both delete them and credibly commit to consumers that there are no more fake reviews and that you should fully trust reviews, they are sort of in the trap we're just deleting fake reviews would make them worse off unless they are able to convince people that they are all gone.

actually out of time. So I'll All just leave it there, and say thanks. right, so I'm Great, now we have Joel wafog to discuss

[>>] OK thanks, thanks very much. And thank you for including me. And it's really great conference, and I really enjoyed reading this paper, how Waldfoegel] does button this work? Green button goes forward, oh goes big The green green The context forward? for this disease reviews are really important button, to

experience goods, I think need to know how important these are, not if you think about you it, I want to put forward the importance of this paper, the ambiguity of this kind of information is really important as the facilitator of net product entry, imagine normally to enter some kind of product, you have to spend an enormous amount of money becoming known to consumers kind of that's a real deterrent to entry, faridkot aggregation nice thing about aggregation or Amazon is make this a little easier for entrance to add product, so we know as proud side, that this pre purchased information deliver substantial welfare benefits to consumers, it's all threatened by fake reviews, now there are a lot of existing papers showing that they exist, Joe's kind of thought of that's interesting, but it still seems to matter.

So I'm not going to worry so much, what I really want to know is what's the wealth or cost of them, which is exactly what this paper goes out for, I was very excited when I actually looked at the paper. All right. So possible mechanisms that they talk about is it distorts the product quality. It underlines consumers compliance and reviews come in and induces firms to compete differently on price since they can't credibly communicate above quality comatose also add to that to that a heavy reliance on old fashioned mechanisms of like product quality like brand, anyways in their view, Five stories, I laughed, I cried, better than cats, it was the best student evaluation.

I ever got like 30 years ago, also uncompensated that's credible and true, it's important and an interesting question,'s ambition ambitious now is a backhanded comment. I don't have 4 handed compliment comments an ambitious analysis to analyze the effect of fakeness on stuff, is very impressive data collection, and then they include the belief measure, and the demand model and explore its impact on stuff for counterfactual was common and the impressive data stuff is really worth talking about. So first of all the daily sales data on Amazon common not just daily sales ranks which many people use, but they also do this in a previous people I guess inferring from the inventory staff just a very cool bit of detective work, they also do other kinds of detective work to figure out which reviews reviewers reviewers are actually fake, it's a super cool cloak and dagger stuff, then, they do a consumer to get a survey a survey rather to get consumers perceptions of fakeness, so that fake review gumshoe work, have graduate students going to these on the dark web, not quite, Facebook, but they find out who all of these many, many fake reviewers are, and then they go and figure out which products use fake reviews and so forth, and that's really amazing, they're able to come up with some statistics, the prevalence that they could use, and I think that's interesting, products purchasing fake reviews the ones in the red have higher scores than their controls, I do wonder just as an aside how can Amazon do this too.

So that maybe this problem come Amazon has a few smart people there as well coming, but we will come back to that, then they listed beliefs about fake reviews, basically, they're able to show us functions of the ratings that products have. And the number of reviews that the products have, people are skeptical of products with high ratings, and I think the earlier draft it was high ratings and especially with few reviews. Now it looks more like high ratings come up. But at any event, that might make it harder to identify things just because beliefs about fakeness sorry, correlated with reviews, but anyways, I will come back to that too.

All right, then what they do they is incorporate the perceived quality variable into the demand model, and they find it as Brett says they find higher price delivers lower sales to law of demand still holes, has not quite significant, and then higher perceived quality delivers higher sales, so again, that's an encouraging result, again, that said, this I will maybe slow down a little bit and just try to say where I was a little bugged on behalf of the authors because all this great work. I was a little bummed that all of this great

worth is a little surprised, I'm also a little surprised because every time I've used something like Amazon data or hourly Amazon data in the context where I can get it, price coefficient is extremely precisely estimating, probably biased, but extremely precisely estimated and also, I think in the studies that use star ratings on the right hand side, that tends to work pretty well as well, so they continue valiantly, but I'm a little hung up.

So I want to sort of talk about some of this, so two ways to perceive One way is to think of the whole thing as a theory guided collaboration exercise showing large possible effects, and then you probably need to do a bunch of sensitivity, to show what stuff remains big and overall, it's not a crazy way to go, can do a good job showing plausible magnitude as a fake reviews on magnitude thermal is the question,, so that's one way to go. And I think, frankly, the compliments enough substitutes is to find more in the sense of addition additional direct evidence comes there's this derivative the derivative of quantity with respect to stars, which is in a lot of people's papers, so, the cool new data among other things here imagine creating your product where there's a lot of favorite news demo where there's a lot of previews perceived to be fake, I'd like to know how that derivative you know the coefficient on stars, the quantity with respect to stars differs across those kinds of context as a descriptive matter 'cause they guide towards understanding why the estimates are doing what they are doing, moreover, prices play a really important role here, and I would lose my union card as an IO guy if I didn't pledge allegiance to the possibility of ratings changing the pricing, bye.

Would also like to see direct evidence of that, that is the fakeness of ratings, not entirely can't believe that, but why not just go ahead and explore that thing directly through render sort of more plausible A mechanism that's an important part of of the story the model, so, some things that I would, some lowbrow things that I would love to see, estimate the models with stars directly, estimate the models with the sales rank as opposed to the better in some sense better quantity data just because it's easier to get more sales rank and again, just to see do things work the way that we expect them to work, why are the coefficients not coming up significant in the main specification of evil, I'm actually probably going to end ahead of schedule, I really want to say this is a terrific research agenda and paper coming.

It asks a really important question, does a bunch of really useful and impressive things to get the data that's relevant to answering it, I think, I'm going to watch this paper this research program eagerly, and I think we all should, so thank you for the opportunity to check it out. Hi, Ben kastner FTC, it's really cool paper, I really enjoyed listening to it, I think there is a couple of areas that are kind of missing, so, on the consumer

[>>] side you are kind of assuming passive beliefs about the level of fake reviews, so, I'm pointing people a lot to the platform governance literature. This there's some papers looking at Catholic, consumers engaging in screening Thomas you could go to websites like reviews meta, or fake spot, and so, on the one hand that is going to reduce the misinformation cost, but it's also going to create an additional search cost or screening cost for consumers, so, the kind of mitigate the benefits I've just missed pressed that you are getting at, also, on the seller's side, there's a cool job market paper to this 2019 I think, you can discuss that more offline, that was talking about incentives to buy fake reviews, and one of the things the author to that paper pointed out, is that this is creating a cost for Sellers, so you're kind of having an advertising style prisoner's dilemma where you are increasing costs that eventually get passed through to consumers, which is also going to have a welfare impact great suggestions, we have thought about Hollenbeck] so in our survey experiment about Great, beliefs come to show people the average rating

and the number of reviews, we thought about secondary 1 where we show them a text or three weeks and see if that helps them, in which case, that would be sort of a direct evidence that people can sort of pay a cost better and get measure of information, and it could be I think plausible also to just directly trade with corporates the cost of information acquisition, and then, on the seller side, yeah, so we don't model the supply side, we don't model the choice to buy fake reviews, we sort of take it as given, but it would be interesting to think about that choice and think about the prisoner's dilemma, and we have a pretty good sense of the cost of fake fake reviews, since it's the cost of the product plus taxes and and fees, so, that could be something that we could explore Hi, my name is Sammy I'm from Georgetown, and partially my question was very closely related to second question,, essentially, if you are a sailor and you are deciding whether you actually want to proceed

[>>] with fake reviews assuming you have no conscience, then comment just becomes a cost of getting people actually on board with this, and I was just curious if you had any estimates essentially, yes, there's obviously kind of you're basically giving them a free product, but in addition to that, is it hard to get a lot of fake does it take to do reviews? that? And in the end essentially, if How you're a company. Even if you have no conscience is it even worth much doing this, how effective time Yeah, we can is definitely speak to some of these things maybe not the household cost come over the sellers for the most part it?

the sellers for the most part are concentrated to a particular Hollenbeck] region in China. I'm sure I'm trading trade secrets and I'm sure it's relatively easy to do, one thing that we can do we explore doing a little bit comments calculating the marginal profit impact of a single fake review, by looking at the extent to which demand changes comparing that to the cost, I think that could help answer that Mark Reisman Boston university, so the results on positive reviews. But I thought there was also a whole world of negative reviews and people being skeptical of the negative reviews.

And I was wondering if you knew

[>>] anything about that in the survey, In the world of fake reviews negative fake reviews. They are a huge deal, not in this setting though. But we don't observe buying fake reviews in these Facebook using we have a sketched out model of why, which is essentially Hollenbeck] that the cost to buying a positive review for yourself is basically just the production cost of the product because you make the sale Amazon gives you the money back, so it's actually pretty cheap, to get a negative fake review to the competitor you have to pay the full cost of the the product plus benefits coming back to you are going to be more dispersed assuming there's more competitors, so it doesn't seem to be taking place in this particular setting, and it seems possible that we're not just missing it, but in other settings, or they don't have to go through this whole verified process Okay, yeah, might distrust the negative reviews even if they they are all real, it .

kind of gets into the side, features of your paper consumer

[Brett the back and see that, products have Hollenbeck] On all sorts of negative reviews are distrustful if

[>>] So they have too many one star reviews. we can Hi Jake Kramer university go of Maryland, so I guess my question is sort of motivated by your context, so I guess I'm wondering how is the impact of fake reviews affected by advertising since you are looking at Amazon specifically, you know,

[>>] fake reviews can obviously sort of boost products in an organic search, but I guess I'm wondering how that's affected by the presence of so many advertisers sell products on Amazon Yeah, this is an interesting question coming in a sense, I think fake reviews on advertising are substitutes, One of the puzzles to me about this is that Amazon makes a lot of money from advertising, they don't make money,

well, maybe they make a little bit of money [Hollenbeck] from paper views, and you know, it seems like they would have a strong incentive to get rid of the fake reviews to the extent that it would cause sellers to substitute into buying more advertising, and yet, the sort of obvious thing that we note is that they are not doing that, so they are allowing still massive emails but fake reviews, and so maybe that says something about the direct substitute ability between the two, or maybe that they are advertising economists, and they are machine learning people who are in charge of fake reviews are just not talking to each other, but this is a very interesting question, question Alright since this is the last last question I'll give two questions, so first I'm wondering if Amazon has a profit right, margin in some of these All markets you can guess [>>] that fake reviews can shift to different brands they may not want to [>>] stop bakery views to shift to their private label, and 2nd I'm wondering if you can assess pink spot and review melody screening tools, because no one quite knows how well they work, and you have some objective evidence on that

[Brett Hollenbeck] Yeah. OK. We haven't looked at the presence of Amazon private labels in the subcategories. But that would be an interesting thing to check, I will say in the survey we sometimes see people on Amazon gift card. And hopefully they say 0% chance this is buying fake reviews and if they don't, and we eliminate them, but we should look and see yeah, because that would also speak to the platform's incentives, and we could directly measure I guess that substitution probably and incorporate that into that measurement of platform profits, and then, I'm forgetting what the second question was, oh yeah, we haven't looked at this directly, so fake spot actually, if you want a sort of fresh rating of fake spot for a product, you have to ask it for it, otherwise, it's the last time it someone would asked it, be so quite time consuming to get fake spot ratings for all of these products, by sort of casual empiricism that that it used to be quite good, and I would check it a lot, and it correlated very closely with our data, now, sort of two years after are projecting it's broken down entirely.

And so that doesn't seem very good, so we can explore that it would definitely be interesting, and it would speak to the sort of transaction cost the cost of information inquisition it currently seems like it's not that helpful Great, thanks Great, well that concludes the 16th annual FTC conference, so we hope to see you at the 17th next year, probably probably here, about the much.

[>>] same time, thank you Just quick logistical thing if you want to leave both your lanyards and your all name because we reuse the cases on the table where you got for them, that would be great,, and thanks so much coming.. .