



Seventeenth Annual Microeconomics Conference | FTC | November 5, 2024

Samuel Kleiner: Okay. Welcome back, everybody, to the second day of this conference. We're going to open today with some remarks from Steve Berry. Dr. Berry is the David Swensen professor in the economics department at Yale University. (We're E.1s)-3.998 M

Fellow of the Industrial Organization Society. At Yale, he has served as the department chair, the director of the Division of Social Sciences, and as the inaugural director of the Tobin Center at Yale, a research center focusing on domestic economic policy, and they're also a sponsor of this conference. He's served as a consultant for governments, policy institutes and the private sector, focusing on questions of antitrust as we.004 (, pt/-dB TJ ET Q q 0 0 612 792 re W* n BT /

Aviv was addressing yesterday in the particular, that I would like to address a little more in the general, which is how do we maintain credibility in a world that is skeptical of us, of experts, of many different things?

Let me start just a little bit with the Tobin Center and what we do. I think we're a little bit different than a lot of other university centers. We were set up

approximately seven years ago, though we've probably only been a staffed enterprise for more like five. And the idea is that so much academic research gets bottled up in journals. One of your junior colleagues writes a great paper, and you read the policy conclusion, and it sounds great, and it goes in the great journal, and they get tenure. But as far as the world is concerned, it gets filed away in the library, and that's the end of that.

There are a lot of centers that hold conferences, and put out white papers and that sort of thing, and I think that does have some effect on the world, but I don't know how much the stack of white papers, or other links on the website, I don't actually know how much effect that has on the world. I think a lot of centers, if you really dig into what they're doing, and I don't mind this because I'm at a university and we need money, . But I think a lot of them are building publicity and donations, and that that's their actual goal when they're holding these policy conferences, and so forth.

We decided to ask the question of what would you do if you really said, no, we sincerely want to make the world a better place through economic research? How would you design a center to do that? It's kind of corny, but I said, what if we were just really sincere? What if we just really wanted to do that? And so, that was our goal, and there's a problem because it's nothing I've particularly ever done. One of the things we had to do was to hire a really good staff, and to talk to the few people at the university, I think, who already knew how to do that. One of those is Zack Cooper, who you're going to hear from later, who was already pretty good at this.

There are some kind of mundane things you have to do in terms of supporting research. If I look at my young colleagues today, I was taught that you take econometrics because economic data sets are terrible, and you have to do a lot of fancy stuff to make up for how bad they are. The younger generation has a better idea, which is just to combine 10 different data sets and get really better data. That's actually a better idea. You still may need some econometrics at the end, but it's much more solid. You need the infrastructure to bring data in, you need to help the university and the researchers figure out how to sign a data use agreement with a government agency, or with a private sector firm or something like that.

And again, these are things that the faculty aren't taught to do in graduate school. You're not taught how to negotiate a data use agreement. You're not taught how to explain to the university's lawyer why your confidential data set is actually secure according to the standard, blah, blah, blah. As an IO economist, one of the things I feel, is we can bear a lot of the fixed costs. How do you appoint a preloc research assistant? What title? How do you run it, provide some support for that program? How do you help get data on campus? How do you help with data use agreements? We can do some of that.

But I think there are things you need to do. One of the things, and this conference is an example of it, is you need to connect researchers with the

actual people who know what the questions are. I was a little skeptical when we started about whether our faculty at Yale, if anything, I think academics are always a little pointy-headed and so forth, but I think our faculty in particular was famous for being methodological, econometrics, theory, blah, blah, blah. It's really interesting how much people's eyes light up if you go to them and you say, there's a related topic to the one that you work on, but people in the world would be much more interested in it. I thought maybe the answer would be, who cares? I do what I do. The more common answer is, really, can you tell me more about that? Because actually if they could do the research and affect the world, that seems like a way more fun thing to do. That's really worked.

I think this kind of conference, it would be a mistake, I think, just to think of it as academics come in, and share research ideas and so forth. I think there's a really important role for a conference like this to help the academic community to figure out what the issues are. I think that's one role of the scientific community, is to pick topics and papers that are relevant to the work, that's policy relevant and not just... some of the papers are just cool research. That's fine too. But what are the interesting topics that academics should know about? I think that's an important role for this conference. From the inside, what are the things we don't know that we should do?

Now, let's say that someone comes to us and they says, I do have this great policy result, and I would like it to influence the world. Well, one thing you have to be careful about is that this is really ready for prime time. There are a lot of research papers that are good research papers, but are not really ready to influence policy. Everybody loves to write a little policy section at the end, but you have to be a little hard-headed about should there be two more papers, or three more papers before we take this out to the world? Or is this really truly credible research that we're not going to feel bad if someone takes this advice, because we don't actually know if it's right? I think that's the first thing.

But the other thing you have to have to do, I think, is to figure out are people ready to listen? Is it ripe? You can get a press release or something like that, you can hold a conference or something, but you're not actually going to affect the world if no one's willing to listen. In practice, I think this means avoiding things that are overly partisan. If things are polarized already, there's just a group of people who want to hear evidence on this side, you're not going to change the debate. Now, the good news for a lot of economists, particularly for microeconomists, is a lot of the stuff we do is a little bit technocratic. We were

was famous for being methodological, econometrics, theory, blah, blah, blah. It's really interesting how much people's eyes light up if you go to them and you say, there's a related topic to the one that you work on, but people in the world would be much more interested in it. I thought maybe the answer would be, who cares? I do what I do. The more common answer is, really, can you tell me more about that? Because actually if they could do the research and affect the world, that seems like a way more fun thing to do. That's really worked.

I think this kind of conference, it would be a mistake, I think, just to think of it as academics come in, and share research ideas and so forth. I think there's a really important role for a conference like this to help the academic community to figure out what the issues are. I think that's one role of the scientific community, is to pick topics and papers that are relevant to the work, that's policy relevant and not just... some of the papers are just cool research. That's fine too. But what are the interesting topics that academics should know about? I think that's an important role for this conference. From the inside, what are the things we don't know that we should do?

Now, let's say that someone comes to us and they says, I do have this great policy result, and I would like it to influence the world. Well, one thing you have to be careful about is that this is really ready for prime time. There are a lot of research papers that are good research papers, but are not really ready to influence policy. Everybody loves to write a little policy section at the end, but you have to be a little hard-headed about should there be two more papers, or three more papers before we take this out to the world? Or is this really truly credible research that we're not going to feel bad if someone takes this advice, because we don't actually know if it's right? I think that's the first thing.

environment or do they get the Amazon cut down? They're supposed to come up with a number, which is how good they are for the environment, and what did they know about this at Treasury, three months ago? They have to make a decision. There are just cases where they're really looking for help.

These are the things that are often the most ripe. Not some sense the big thing, like should taxes go up in general, or should taxes go down in general, or are we going to revamp the tax system around tariffs, or carbon taxes or something? The academics can talk about that, which is valuable, but we're not going to affect things in the short run.

Then the other thing we need is we need people who are good at that. So again, Zack is good at that, but that's not primarily me. I think what's been interesting is the degree to which we've been able to get just fantastic staff who actually want to get out of the partisan environment, and actually be able to talk to both sides and not be trying to score a points. It's interesting to me how surprised people are that there are people in Washington DC who are sincere in the same sense that I meant sincerity, that they actually want to do good, they actually want to make the world a better place. And of course, that's one of the inspiring things about this conference, is there are a number of places in the federalame

ago, or they'll be asked it again. And I think by getting that out to a broader set of policy-oriented economists, and these papers are not typically going to... they're going to be solid research. They're not going to be the policy advocacy papers. But other people can take them and see what the policy implications are, and help get those out to the world. Maybe when, in your ~~research~~ research life, you're focusing on the question that the people are asking you.

I think what I want to finish with is just, it gets to what Aviv said yesterday, which is that I think in all of this, we really have to think about our credibility. It's going to be super important that we're always giving, doing that first step of figuring out what's the good advice of being honest when we don't know so much. We sponsored a whole big digital markets project, and I said I want to be careful how we market this, because this is super new stuff. I want this labeled as, these are the thoughts of very smart people who have thought about this a lot, and are talking about policy. It's not like I can sign off and say, yeah, every one of these statements they're making, we know are correct yet. And that's fine, and a policymaker may say, well okay, but I got to do something, and that's the best we got. But we should be honest, that's not a finished research agenda.

Whereas in other cases, we just help fund a paper on the effect ~~of~~ ~~one~~ mothers' labor force participation. Done right, it's a huge effect, and I just really believe that result, it's a super careful paper. I just think we can go up and say, look, this is a big result. It's a big, new result, and I really believe that result

As a preview of our results, obviously the harm to workers depends on the configuration you're looking at. We find that most of the harm happens when there is both labor and product market overlap. Those seem to be sort of additive in terms of how they affect workers. And workers can actually benefit in some cases from the mergers, when there's only product market overlap. That increase in bargaining surplus mechanism can in some cases dominate, and actually increase wages for workers.

We also find that the Delta HHI, or the change in concentration, is very predictive of both worker effects and consumer effects from mergers, so that's very helpful. And we also find that conventional merger simulation that's based

other products owned by that firm. And so that's going to be the consumer substitution, the recapture, that allows the multiproduct firm to actually have a higher threat point. This is how that third mechanism I was talking about, the bargaining leverage of the firm comes in. As that set ZJ gets larger, the firm's threat point is going to go up.

negotiations and also the downstream price game. But in equilibrium, the wages and prices satisfy all of the first order conditions, so there is a lot of interrelationship between the bargaining and the downstream game. I'll go into a bit more detail on the methodology that we have for the simulations now. The data that we need for the downstream for [inaudible 00:33:11] conditions is pretty standard. We just need margins, costs, market shares to get the downstream demand parameters. The bargaining for [inaudible 00:33:19] condition. Additionally, we're going to need data on wages and on the workers' outside option wages to identify the bargaining power parameter. We could alternatively have assumed a bargaining power parameter and then backed out wages or outside option wages, but we chose to estimate the bargaining power parameter.

These are pretty standard things that we would receive during an investigation at the FTC or DOJ, but it's more difficult to find public data at the firm level, margins and wages. And especially for our purposes, we wanted to do a broad set of simulations to really illustrate the model. So we needed industry ride for margins and wages. This was difficult to find. Definitely talk to me if you know of any other industries we could look at, but we did find a few applications that we could use the data for.

The first of those is the US hospital industry. So building off the work of Ellie Prager and Prager and Schmitt, we're taking the HCRIS data to get wages at the hospital level, the CMS data for some cost and price information, AHA for ownership and location, and the BLS for outside option wages. And we're going to, again, focus on the market for nurses and pharmacists. So following Prager and Schmitt, we think they might have specific human capital and therefore be more likely to be subject to mergers and market power.

Another advantage of taking the hospital data is that there are some off shelf market definitions that we can use, which is helpful to us since we want to do a lot of simulations and don't necessarily want to do a market definition exercise for each market. So we use an HSA, hospital service area, as a narrow stylized market. This is a zip in which residents receive most of their care.

We also take off the shelf the HRR, which is like an amalgamation of HSAs. You can think of a broader market that's made up of multiple HSAs and we use that as our broader labor market for the second configuration.

The second data that we use is the Colombian Manufacturing data, 1991 census. This is a lot of different industries in Colombia and a lot of different regions. So we're going to

So just to recall, the four mechanisms that we're trying to illustrate here are the downstream effects on employment through the standard product market simulation, the increase in bargaining surplus, and then the changes in leverage that tend to decrease wages.

To really get at those four mechanisms and disentangle them, we're actually going to run three different versions of our model and compare them. The first is just your standard downstream product market merger simulation. The second is a sort of modified bargaining game where we're only going to allow the surplus to change and the firm's leverage to change. This is very close to what Horn and Wolinsky did. And then finally we'll do our full simulation, which adds in importantly that fourth mechanism, which is the decrease in the workers' leverage.

We run all of these mechanisms across the three configurations I described earlier. So for the case where there's a change in both product and labor market surplus, we're going to take in the hospital industry HSAs, so very narrow geographic markets for both labor and product market. For the case in which there's only labor overlap, we're going to take that broader HRR in the hospital industry as the labor market and the narrow HSA as the product market. So in this case, there would be no change in the product market HHI and only a change in the labor market HHI.

And finally, in the third configuration, we take the Colombian Manufacturing data and we'll look at mergers 0.00000912 0 61R(ery n30.72 Tm 0 g 0 G [(of)] TJ ET

potentially pretty important. Moving on to our next application, we look at labor overlap only. We have 324 mergers here. By construction, there's no change in the product market HHI here, so the downstream only simulation is not going to have any effect on workers and the Horn and Wolinsky simulation also. But there is a change in the Delta HHI in the labor market. And so in the full model we do see a pretty big decrease in workers surplus. Again, pointing to the fact that labor market overlap could be pretty important in the case when there is labor overlap.

Moving on to our third configuration, we have 423 mergers in the Columbian Manufacturing data. Here in the downstream only Horn and Wolinsky model, we see a decrease in worker surplus. When we move to the full model, in this case, it's a little interesting because there's actually no removal of an option for workers since there's only product market overlap here, no labor market

And another sense, this is actually really surprising because we're seeing that wages are going down pretty starkly here, and yet we're not seeing pass through to price or output. We do see the direction of the correlation there is in the correct direction. So price is going down very slightly as we get to higher Delta HHIs and output is going up very slightly. But the reason that we don't see more passthrough in this case is actually because of using the hospital industry.

So the share of costs of nurses and pharmacists of the hospital is only about 2%. And so when you're seeing a 20% decrease in wages on an input that's only 2% of your inputs, it's just not going to be passthrough meaningfully to consumers. But that's sort of an interesting result here. We did do some Monte Carlo simulations early on. If you push the percent of the cost up to more like 90%, you can actually get stronger passthrough.

Oh, I think I missed. Yeah. Okay. So the last case is Columbia Manufacturing. Here again, we see a strong negative correlation between wages and output and the Delta HHI, which here is only in the product market because there's only product market overlap. What's really interesting here is that if you look in the left graph at wages, actually in about 15% of the cases, worker wages are going up as a result of the merger, especially when there's more of a benign merger. So Delta HHI is less than a thousand.

So this is pointing out that that passthrough mechanism, the increase in bargaining surplus, increasing wages, actually can happen and does happen in about 50% of the cases here, although overall workers are generally much worse off, especially when you take into account the output going down, employment going down.

The last set of results I'm going to show is looking at enforcement screens. So trying to see how good downstream product market simulations do in terms of capturing worker welfare. So this is very close to our heart as an agency. Thinking about if we were just blocking product market overlaps, how well would that do in terms of protecting workers.

So the metric we're going to use for this is if we blocked every merger that would cause more than 1% decrease in consumer surplus, how would that do in preventing mergers that would harm workers more than 1% or 5%? And obviously this is based on the calibrations that we did, but here are the results across our three configurations.

In the case of both product and labor overlap, if we blocked all the consumer surplus decreasing mergers, we would catch 77% of those mergers that harm workers more than 1% and 99% of those mergers that would harm workers more than 5%.

And the results when there's only product overlap are also pretty good. 45% of the mergers that would harm workers more than 1% and very close to a

hundred percent of the mergers that would harm workers more than 5%. So that's pretty good results. The key thing is that if there's no product overlap, by construction, this metric cannot help us here. And so we're catching 0% of the cases where there's only labor overlap.

So I'll just conclude now. We did a two-level vertical supply chain simulation. We calibrated it to two industries, hospitals and manufacturing. We found that workers were most harmed when there's both product and labor market overlap, but that workers can actually benefit in some cases when there's only product market overlap from that pass-through of profits.

We found that the Delta HHI is very predictive of outcomes for both workers and consumers. Initially we looked at the conventional product market screening tools and found that they're actually pretty effective. The only key case that they miss is cases where there's no product overlap and only labor market overlap that'll be most at risk for error.

That's all I have. Thank you so much for having me.

Speaker 1: To discuss that paper is Elena Prager from the University of Rochester and the MBER.

Elena Prager: Good morning. What's the

Related to that, I think there are some other things we want to be careful about if we're going to try to use HSAs for both product and labor market definitions. I came of age as a health economist and so I am sorry I can't help but put up a

agreement. So here, how should we think about when the firm chooses a quantity and how it's responding to wages and how should we think about who exactly the labor side is caring about with an agreement and without an agreement?

Miriam LarsonK...: Yeah, it's definitely a really difficult choice. I think we made the same choice as Alan did in his talk yesterday of having the union cares about wages times output. So it's kind of both, but that's a bit of a cop out.

And then the firm is just setting the quantity according to the product demand. So we don't have any kind of labor supply elasticity that's constraining the firm in terms of the number of workers it can hire, but maybe we can talk after if you want to give us some citations to that older literature.

Speaker 2: All right, and one last question.

Audience Questi...: I had two comments. First, excellent paper. I enjoyed it. First following up on the comment that was made. When you're bargaining, you're bargaining just over W , but you can be bargaining over W , setting W and L so that it's... You're not on a demand or a supply curve.

But the second question I have is in your simulations, I think you could do the experiment, the actual experiment of what efficiencies you need in order to offset these harms. And does an efficiency have a bigger effect in the product market or the labor market, and under what circumstances would it? What's the minimum efficiency you need in order to make the merger desirable?

Miriam LarsonK...: Yeah, that's a really good suggestion. Yeah. There could be interesting interactions of the efficiencies in both markets. So yeah, that's a great idea. And in terms of bargaining over W and Q , I think we'll have to give that some thought because it's hard to reconcile the product market demand for Q versus the bargaining over Q . Maybe I'll talk to Alan later about how he figured that out.

Speaker 2: Thank you.

Speaker 1: And now for the second paper, Benjamin Vatter from MIT.

Benjamin Vatter: That's a fancy clicker. Let me see. Right? Do I see myself anywhere? There you go. All right. Where are my slides? There you go. All right.

Thank you very much for including me in the conference. I'm very excited to talk to you about this work on vertical integrations. So the motivation for the paper is broadly that the role for vertical integration in healthcare market has been steadily growing. Around 80% of our physicians are now integrated with other hospitals or insurers. 70% of our drug coverage is now integrated between our PBMs and insurers and around 50% of all inpatient care in this country is now

delivered by hospital systems that are also in the business of selling health insurance.

Now this paper is concerned with this latter type of integration, the one between insurers and hospitals and its welfare effect. And this speaks to

PART 2 OF 5 ENDS [01:04:04]

Benjamin Vatter: of integration, the one between insurers and hospitals and its welfare effect. And this speaks to a long-standing literature on vertical integrations, on differentiated product markets and on integrated care. And both of these literatures gives us pretty ambiguous things to start from in terms of the welfare effect. On the positive side, they say, "Well, eliminating double marginalization is likely a good thing." It's fundamentally a friction on the market. It might increase coordination of care. It might eliminate incentives for hospitals to engage in wasteful examinations that provide very little value and the only reason why they do it is because they're not facing the true price of it. But it might also consolidate market power upstream and downstream. It might lead to incentives to foreclose rivals from access to valuable hospitals and reduce access to care overall."

Now despite the interest in this kind of work, there's been limited empirical research on this because large impediments on data and setting. On the data side, it is very hard to track ownership in a lot of these settings. And people in this room have done enormous amount of work to track, for example, ownership of physician groups and hospitals. And on the setting side, what happens with vertical integration in particular is that often they come with large organizational restructurings and you need to give them time to really show up in the data. And in a market that is very active, like the American hospital market, by the time you look at it again, you know there's been a lot of other horizontal mergers of hospitals and it's not clear what you're looking at.

So we're interested in this question and because of these challenges, we're going to go and look for it in another setting. We're going to look for it in the context of the Chilean healthcare market. And in a couple of slides, I'll try to help you map the Chilean setting to the American setting to the best that I can, but fundamentally, we're going to look at a privatized healthcare market where there's private insurers and private hospitals and contract structures that look similar to what we have in the US. And it has a richness in which we can observe the contract struck between

there is five private insurance. And I should say this is the healthcare market. There's a bazillion regulations around it. I'm happy to talk about your favorite aspect of healthcare regulation offline. Here's the key things you need to understand for today's talk. And so there are five private insurers. They're offering individual or family plans. It's decentralized. You go to a website and you enroll. And the regulation says that there are different age groups and gender and family comprising. So basically if you're a woman between and 34 with no dependents, you see one set of plans and premiums. And if you're a man, 35 to 45 and you have dependents, you see another set.

In this market, there's negligible deductibles. There are no out-of-pocket. Really for inpatient care, the most relevant thing, or fundamentally, the only relevant, is this coverage rate, which is one minus the coinsurance, which is a fraction of the bill that the insurer will pay if you go to a hospital.

Inpatient care hospital has to be a base tier or preferred tier. The base tier (the best) is around 55% of the bill. And at every plan, every insurer (the decision) to have a network of hospital to call preferred and to have a preferred coverage, so 88% of the bill. And that effectively on the grounds a lot like having in-network and out-of-network providers. And that's going to be a negotiation of what we would have in-network providers here, so it's like the preferred tier hospitals, and that every insurer can create various structures of this kind of in-network, or preferred and non-preferred.

On the hospital side, there's 11 large inpatient hospitals providing our services.

really matters because our upstream model, our upstream providers here, are heterogeneously distributed horizontally differentiated in meaningful ways. So some of them are best at providing tertiary care. Others are good at infectious diseases. And others are good at oncological care. And you as a single patient at risk for all of these types of conditions of treatment are better delivered at different types of hospitals.

provide almost uniformly quality across the board. And I'd think about them as those are the ones that are benefiting the most from getting excluded from isolated from high-quality integrated hospitals.

There's an outside option here both on the public side for insurance and hospital care. It's a relevant outside option for the hospital side. The insurance market is really segmented, so you can understand almost everything here just thinking about private end release on their own. On integration, two of these insurers are integrated. They own three hospitals, one is a hospital insurer pair that integrated through our data and then stops being integrated into a fight around surplus splitting, and then they stop being integrated. We leverage very minimally, and I'll how

exactly what for. But to give you an idea, VI insurers accounts for 60% of admissions at their own hospitals. They're very, very successful at steering people to their own hospitals, and they account for only 13% at rival hospitals.

And they do so to a large degree by being more generous. So a VI hospital is 33% point more likely to be preferential at its own insurance plans than at rival insurance plans. But to be clear, there's still going to be integrated insurers offering plans that have preferential hospitals that are integrated with rival insurers because those VI hospitals might be the best option for, say, maternity care at downtown Santiago.

So here's our model. And the first stage insurers are going to design their plans. They're going to design their networks and their coverage. And then they're going to go and negotiate prices with insurers and hospitals, and insurers are going to simultaneously set premiums. Then consumers are going to see that, going to decide on what to enroll in, and then if they get ill, they're going to decide where to seek care. So stage two, three and four are fundamentally a [inaudible 01:16:08] town model, a Ho and Lee model, with some adjustments to the Chilean setting. And really stage one is where we bring something new to the table.

So if they get ill and they're enrolled in a certain plan, a consumer, I, that has some diagnosis, D, is going to choose to go to a certain hospital based on how much they have to pay out of pocket, what is it distance that they have to travel. They might be influenced by VI marketing, and this is what we leverage this integration at. VI insurers are very, very good at marketing their own hospitals. There's systems by which you can call and ask, "Where should I go for these kind of things?" and they will very likely to tell you, "Well, go to our own hospital." And we can see... This is what we can see disappearing when the disintegration happens. And then there's hospital diagnostic [inaudible 01:16:52] effect which captures perceptions about this horizontally differentiated quality, which largely matches what we know to be true in the market. So people seem to understand that certain hospitals are better at

differentiated networks. So you should expect even an inefficient outcome here, not everyone picks up the same insurance plan.

Now on the insurance side, insurers are going to set premiums to maximize their own insurance profits, and if they're integrated, some hospital profits. The insurance profits are very simple. It's just the likelihood in which every consumer picks up my plan times the revenue that I collect from it minus cost. There is no risk adjustment in this market. In the paper we document, there's adverse selection, there's underprovision of private insurance as our theories predict it is there, but the market still subsists. And then if they're integrated, they're going to care about their hospital profits. Following the literature on vertical integration, we're going to allow for that kind of internalization to be imperfect or more than perfect in a way. We're going to identify and estimate a parameter. It is close to one. For the sake of today, you can just think about it. This data is being one for everyone that is integrated.

Integrated hospitals and insurers, they're setting exactly the same objective. They're solving... They're just maximizing over else. They're maximizing over price. That's the only thing that changed from here to here, maximizing over price. And if they're not, they're going to bargain a la Nash. And this negotiation is very standard. If a hospital system disagrees with an insurer, the whole system gets disconnected. There's passive beliefs about what happens in disagreement about what everyone else is doing. The only thing that is Chilean specific here is that the law stipulates some penalties to insurers for violating past access. And so there's some legal penalties for disagreements that are happening here. They still happen, but there's some right for consumers to sue and get some benefits out of that. It just explains some rogue agreements that you wouldn't explain otherwise. And the interesting part, the new part, is on this stage in which insurers are going to decide for every plan what is the

tiering. And actually Ellie Prager in our job market paper already told us that when network are tiered, it's easier for people to shop around. And we see that in practice. But in Chile, you can actually call a hospital and say, "Hey. How much would it cost me?" And they will tell you a number. And that's in part because the system is simple that way, that they're not going to ask you, "Which number plan are you in?" It's like they have one price, they have one tier and they can tell you. And so that helps people shop around. And there are difficulties. You're still shopping for something that's highly uncertain about the whole package of services that you get. That is of course difficult.

But an interesting thing about this to note is that because people value premiums way more than... Respond to premiums way more than prices, the rising rivals' cost channel is going to be weak because the downstream insurer has an advantage here. When the hospital increases its price by \$1, it can offset it by less than \$1 decrease on the premiums. People value the dollar in premiums more than the value in cost. And in equilibrium, and it's important note that this is an equilibrium outcome because it's in fact determined by the networks that they're selling, they have this advantage and they can offset that increasing cost by less than the equivalent decrease in premiums. And so the rising rivals' cost is weak.

Now, if you're concerned about what is the role of this elasticity that you might not be familiar with on the prices, the main thing that it's doing for us is that it's preserving this behavior of in and out of network. It's keeping people within network. This is what it's doing. And even though in the US we're quite inelastic on the intensive margin of care, we're very elastic to moving out of network, on the extensive margin. This is what this is doing for us primarily.

On the bargaining side, we estimate moderate bargaining weights, which is first of all, what it's telling you is that model is able to rationalize the advantageous position in which vertically integrated systems are, not through bargaining weights, but through actual market power and value that they provide in the system. And we have a formal results that tells us that in this relatively common bargaining framework, if we did not have auxiliary data on hospital costs, which is not the full cost, but it tells us a component of the cost, our system would not be identified. We would not be able to jointly identified the remaining components of hospital costs and the bargaining weight simultaneously. And this is, I think, a general result on cross-sectional identification of these bargaining models.

And so we take this. We go into 2016. It's a later part of our data. And we ask, "Well, we have an environment with vertical integration. What would happen without it?" We break the vertical integration linkages and we let insurers reoptimize their coverage, their plans. We let hospitals and insurers renegotiate their rates, insurers to reoptimize their premiums, and people to choose again what they're buying. Starting from the plan design, we can see that... So here what I'm showing you is the base coverage and the preferential coverage is what's highlighted. So for example, for a VI insurer, on average, the base

different kinds of care. Insurers are financial institutions. If you take away the integration, they're fundamentally identical. So competition among them is way stronger after you ban vertical integration. And the countervailing power here means that hospitals are in much better position once you eliminate this differentiation that they had through vertical integration downstream. They're in a much better position to substitute one insurer for the other, and that gives them an advantage in their negotiation.

So the average rate, the average price there, increases, but the average price that consumers experience does not because the ~~VI~~ insurers are doing exactly the same play as VI insurers. They're steering consumers away from where this increase hurts the most. They're taking consumers back ~~for~~ ^{above} care to downtown Santiago for ~~low~~ ^{low} complexity care at cheaper hospitals. And they're moving high complexity care or retaining it at the Star hospitals that now face increased competition. So there's a lot of ~~sorting~~ ^{sorting} of patients specifically away from middle quality that seems to be largely distorted by the incentives created by vertical integration.

If we look at the demand, so this is 2016, so it's a bit different than the 60% number that I showed you, told you before, if you look at a VI hospital, 72% of its admissions, which you see in the baseline, is coming from its own insurer in our baseline 2016 model. 24% comes from other ~~VI~~ insurers and a little sliver there, 4%, is coming from the other VI insurer. But when VI is banned, this basically equilibrates, which telling you that insurers become essentially identical. So now the admission rate at a VI hospital and ~~VI~~ ^{non} hospital look essentially the same across these different insurers. And this speaks to the fact that access has become basically uniform. Those diagonal lines go back to being essentially the same.

So who benefits? Well, the VI firms lose around a qu 0.00gg.68 4c

rivals' cost channel was already weak, and so you're mostly introducing double marginalization, which is fundamentally just a friction.

Now, that's all good, but this is an unstable situation because you're creating a hold-up problem. In this world in which you leave plants as they are, you have vertically integrated or formerly vertically integrated insurers, stuck with plans that the only reason to exist is because they were trying to steer patients toward their formerly integrated hospitals, which they no longer own. And now those former partners are going to hold them up on that value. They're going to say, "Well, you have a plan. The whole system, your whole business model, relies on me giving that coverage. And I no longer take a share of your surplus as an insurer, so I want to be paid for it." And therefore, it is clear that the insurers are not going to continue to retain those plans and keep offering those plans. And this is why we had to endogenize the plans.

And it's also... To be clear, it is not as simple. If you do this analysis and you're thinking about, well, going to study these kind of situations from an antitrust perspective, the plans that we see in the counterfactuals are not the plans that the rivals were offering. It's not as simple to say, "I'm going to take whatever the rivals are offering, going to put it into what the VI is going to offer and going to do that simulation." That's completely wrong about what happens, because if you remember the plans that the rival is offering, it's fundamentally this other gray line that was also had weak diagonal connections. So fundamentally, breaking vertical integration breaks this VI hospital silo and changes the welfare effects of vertical integration.

So last minute to conclude. So vertical integration distorts equilibrium prices, coverage and access. The price effects of vertical integration in our setting seems to be largely positive. Again, double marginalization, eliminating it is largely a good thing and the rising rivals' cost effect seems to be relatively weak in our setting. But it is through the plan redesign that things flip, and that looks a lot from the static picture like foreclosure, but it comes from a different source. It comes from the downstream decision of changing who you're sourcing. So it harkens back to the foundational literature of VI that thinks about input choices. But here, the fundamental difference if you go back to this literature is the upstream is differentiated, while almost all of our theoretical literature is written on differentiated downstream players. And it creates slightly different mechanisms.

Now, we find an outstanding role for the competition over plan design in healthcare markets. It flips the welfare sign. And we provide a methodology for solving these kinds of problem at scale. And I didn't have time to talk about this today, but we provide robustness and say, "Well, even though in our setting we don't see quality or cost gains from vertical integration, we can give you a number and say, 'Well, how big would it have to be?'" I can tell you that the cost efficiencies that would make vertical integration welfare neutral here are really,

Shoshana Vasser...: ... In H8. That means that these guys are all in one network. That's what we mean by integration. Okay. Now, the bottom line that Ben showed you, I'm just going to reiterate, this is just one piece of the many results that are there, is that vertical integration is bad, and specifically if we were to do the comparison, take vertical integration, take the market as it is and compare it to a counterfactual in which we turned off vertical integration altogether, the net welfare benefit accounting for consumers, accounting for the firms, for the hospitals, accounting for the insurers, accounting everybody's total welfare, we would get a net benefit of about \$40 million or dollars, right Ben? It's not in Chilean Pesos or something? Okay. All right.

So the \$41,710,000 is the figure that's in the bottom there. If we work our way up, first of all, we see as, I highlighted here, that consumers are benefiting and there are a big part of this \$40 million. But another big piece is this thing called moral hazards spending, and this is a little bit different than what we often think of as moral hazard in other settings that I'll talk about it in a little bit. So just keep that in mind.

Okay, so how did we get to this number? So basically there's a large, large, large modeling exercise, and I'm not by any means trying to argue that we don't need this, but I'll walk you through the steps and then decompose them. And so what we do in order to get to this number is we first take a model of consumer demand for hospitals. The reason we can do this in part is because the Chilean insurance system works more like a ~~two~~ part tariff than the way that insurance works in the US in the sense that you pay a premium upfront, and then you basically pay per service when you go to the hospital. That's the reason, I imagine that, as Ben said, when you call up the hospital and ask for a price, they can give you a price. They can give you a price because you are actually going to pay a price that is foreseeable and is not dependent on some sort of later stage bargaining. That's really nice because you then see consumers making choices over hospitals for specific conditions and this is something that's in their data that will allow them to see how consumers trade off prices at different hospitals and presumed anticipated quality of care at these different places given their conditions and given the insurance plans that they're currently in. Now, they take this model of consumer demand, which is already a feat of an estimation exercise, and they put it into a model upstream or downstream I suppose, of consumer demand for insurance plans. So you anticipate that you know what conditions you're going to have and expectation. You make your insurance plan at the beginning of the period as a function of the premium that you anticipate and the utility, the inclusive value of what you're going to get once you're in that plan. So now we have two orders of demand estimation and we feed that into a model of price setting.

This is interesting. It happens at the same time in their model. They basically have the hospitals and insurers bargaining over prices. So they bargain over the split of revenues from the consumer side or the split of the ~~cost~~ ~~cost~~ between the insurer and the consumer for every transaction. And at the same time, they have the insurance companies setting their premium choices. So the

insurer is basically simultaneously choosing all of the prices that it's going to encounter, the revenues it's going to get from the consumers and the amount that it's going to spend paying the hospital for every bit of care.

Once we have that price setting model, we then feed that into another model further up or downstream, however you like to think about it, where before the insurers decide on all the prices, they decide on the design of the actual insurance plans. Specifically, they design how much coverage they're going to assign to every hospital treatment pair, and specifically one important margin is going to be which hospitals they include in their preferential tiers or they give preferential status to in different plans. This is going to turn out to be really important and I'll show you why in a little bit. And once we have all of those different pieces, they then integrate that into an equilibrium simulation to try to figure out how the market will equilibrate, which turns out to be quite complicated, really, mostly because of the insurance plan design, and then they can do their simulations and compare.

Okay, so how did we get there from the perspective of the data? That was the perspective of the model. So I'll show you a couple of snapshots I get at this. So one snapshot, there's no modeling at all in here. This is just raw data. Here I have a pair of tables from some appendix, appendix, E, 351.5 data. Here I

admitted to hospital H2? And you can see that there's a high correlation between these two, the darkly shaded boxes on the left and in the right, and we

Speaker 5: We have time for a few questions.

Speaker 6: Hi, thank you. This is really fascinating and I guess by virtue of being here in the US and US regulatory agency, I'm thinking, how do we apply these results to what's happening here? And I think one of the arguments you hear a lot about here in the US is these are vertically integrated health systems, not just hospitals. And one of the arguments they make is that we're actually keeping people out of the hospital to begin with. And this may be asking too much, if there's any way to respond to that kind of argument within the limits of the kind of data and the setup you have here or just if you have any thoughts on that?

Benjamin Vatter: Absolutely. So these are also integrated health systems and they will tell you the same. You ask them what sort of initiatives do you have in place and they'll have a very hard time to tell you what they have. Now, in practice, we could look at things like readmission rates, there's nothing there. We could look at [inaudible 01:50:32] infection rates, nothing there, right? It's like whatever it is that they're telling you that they're doing, we don't seem to find it. That's not necessarily true for outpatient care, which is not our focus. So there is some evidence that they are able to do those things, the outpatient care. It just seems that managing these kinds of integrated operations on something that is complex, which carries the volume for inpatient care, it doesn't seem to be happening, at least in our data. It doesn't mean that it couldn't. In a way, that the environment also tells you that they might not have the incentive to do it. It's very profitable even without it.

Speaker 7: On your left, here. Okay, very interesting and thought provoking. My question is can we apply what we have learned here to self-referencing for online platforms? Because I'm thinking online platforms can redesign the search algorithm, which is similar to the insurance coverage, but I guess one difference is here, you have the patients to be single-homing, but for online platforms, they are multi-homing to different platforms. I don't know how much of that will change your story?

Benjamin Vatter: Yeah, very good too. I've tried to write the theoretical model behind these things and it looks a lot like a platform model. And in fact, the fundamental is multi-homing. In a way, what this is saying, if the VI system preferences and isolates, does it push away everything that they're not selling away from their platform? Then the other platform that is selling not the integrated products, not the essentials, then those other products on the other platforms are going to face less competition and they'll be able to increase prices. So it's very much within this line. If you write the theoretical model, essentially all of this model, you can write it without any risk, the risk coverage, and this plays a very minimal role. It's fundamentally a model of platforms and heterogeneous inputs. And so I think it's very much within this line.

Speaker 8: Hi Ben, over here.

Benjamin Vatter: Oh, hi.

Speaker 9: Okay, so this paper was huge when I saw it five years ago at NBERs. I'm glad you spent another five years making the paper even bigger. So you've probably done both of the things I'm going to ask about. Thing number one, which is I think an obvious thing to ask about, is are any individual vertical mergers if you were to unwind them, counterfactually welfare enhancing? And then secondly, I think this is probably a pretty good setting to study the idea that maybe vertical integration itself is a strategic complement. If my rival integrates, I want to get siloed and I have a stronger incentive to integrate and there's talk about merger waves and mergers as bank runs or this sort of idea seems like this would be a good opportunity to study that.

Benjamin Vatter: All right, let me start from the second one because that's very interesting. That basically goes back to the foundational theory. The foundational theory on vertical integration actually thinks about endogenous integration. And a lot of what's happening with the star, for example, star hospitals is if you write a model, they wouldn't want to integrate. They just benefit too much from not being integrated, but the history is on your side. There has been new integrated entrants in this market because the incentive is there. Now, it is very much a compliment. So you're absolutely correct on this, it's not part of our paper, but maybe there's a paper to be written with the new entrance on this.

Now, the first part was about... Remind me?

Speaker 9: Individual mergers.

Benjamin Vatter: Individual mergers, right. So there's a very interesting thing. We do this analysis in the paper. We say, well, we have two integrated systems. What if we just kept one? And it turns out there's this thing that, well, if you take the high quality, like the ones that own the highest quality of integrated hospitals, and you just left them around, that's good for the average consumer surplus. It is really bad for 98% of consumers. And so it's hard to tell, right? Because if you look at the average, the average is very misleading relative to the average consumer because it's really, really good for the patients that have this enormous value for going for this really high value specialized care there. But for most consumers, it's bad. And if you just kept the integrated cheap hospitals, which is serving a large demand of very price elastic consumers and you further reduce their prices by eliminating [inaudible 01:57:03], that's good for, I don't know, 50% of consumers, but on average welfare is bad.

And so it is there in the analysis, but it's actually quite tricky and it's an interesting way of the heterogeneous impact of these integrations.

Thomas Koch: So I'm happy to introduce Zach Cooper. Zach Cooper is an Associate Professor of Public Health and Associate Professor of Economics at Yale University. He also serves as the Director of Health Policy at Yale's Tobin Center for Economic Policy. He's a health economist whose work is focused on producing ~~data~~ scholarship that can inform public policy. In his academic work, he's analyzed the impact of competition in hospital and insurance markets, studied the

influence of price transparency on consumer behavior, and explored the causes of surprise out of network bills. We look forward to his discussion.

Zach Cooper:

Hi, everybody. All right, it's down here. Well, thanks everybody for coming for these last two days. And on behalf of Steve, the Tobin Center, all of us, we're so delighted to be able to partner with you and work with you in putting this together. This is, I think as Steve said, exactly what our mission is. So what I want to do is present the first two in a series of papers that we're working on, think about the causes and consequences of rising healthcare prices in the US. And the first paper, which is joint with Stuart Craig and Lev Klarnet, Zarek Brot Goldberg is thinking about whether there's too little antitrust enforcement in the hospital sector. The second paper, which has those guys and also includes Etai Lurie and Corbin Miller from the Treasury is thinking about what are the downstream consequences of rising prices. How does it affect workers outside the healthcare sector? And there are really three takeaways from our work.

So the first is we do think there's evidence of underenforcement. We think that's mostly because of funding constraints. It's not the FTC doesn't know a bad deal when you see it, you do. It's that we think the agency's funding is pretty restricted. The second is that the

And I think part of what's suggestive that it is funding related is when we look at the nature of the deals that are happening across time, the deals that are flagged by the horizontal guidelines and where the FTC is taking action, we see the story begin to play out. The average deal in our sample is raising the HHI by about 460 points. The average deal flagged by the guidelines is raising the HHI by 1800 points. And the average deal where there's been an enforcement action raises one of the merging parties, HHI, it's the maximum increase we see in a particular transaction by about 3,400 points or a predicted price increase of about 23%. And so these are really, really problematic deals. You're finding those, but it turns out there are a bunch of 5% to 15% mergers that appear to be going through. And what we want to do in the second paper is say, "Well, what are the consequences of those mergers occurring?"

So the reason I think it's so important to look at this is because of employer sponsored health insurance in the US. The modal individual gets his or her health insurance through an employer. And what this does is it creates a mechanical link between what's happening in healthcare markets, consolidation, for example, and what's happening in labor markets. And we can go back to some of the theory that Larry Summers sketched out in the late 80s that says, "Look, when the cost of fringe benefits go up, that's paid for by workers." And embedded in most of the theory is this idea that workers value the benefits that they're getting, and if they value it dollar for dollar, in theory, there isn't a distortion. And a lot of the literature has looked at the impact of new benefits. Right? Extending example maternity benefits and found dollar for dollar passthrough into wages.

\$25,000 a year workers, or they can save \$6,000 by letting go of one \$100,000 worker.

So when healthcare costs go up, they much larger proportional share of lower income workers. And so we think they may potentially be the ones who bear the burden of rising prices. And Amy and Owen have done some simulations that actually suggest that if you think of it this way, over the last 30 to 35 years, rising health spending could potentially be the leading driver of income inequality in the U.S. Larger than the effects of trade and outsourcing, automation or a lack of growth in real minimum wages.

So what we want to do in this paper is trace through the causal effect of rising healthcare prices on these downstream labor market outcomes. Now the challenge is you don't want the causality to go right to left. So Nvidia is doing pretty well. Maybe that increases demand among their workers for health insurance. We know from Kate and Robin's work that may increase prices. So we want something that's going left to right in terms of causality. That's where our hospital mergers come in. So we think it's a useful shock to think about the downstream consequences of rising prices. We also think, given the importance of mergers in and of themselves, this allows us to say, "Okay, what are the downstream consequences of these transactions?"

So we're going to bring together all this rich data. We've talked about the merger database and the claims data. We're going to get data on insurance premiums from the Department of Labor for a fairly small set of firms who fully insure via work with Itai and Corbin who are at the Treasury and IRS. We're going to have access via them to the universe of tax returns, securely. I should also say sadly that this doesn't therefore reflect the views of the U.S. treasury because they're there. And we're also going to have access to the CDC's restricted mortality data. And I'll tell you about how that feeds in just a second.

And to give you the sort of punchline empirically, we see the dollar increase in prices raises health spending by a dollar, unsurprising. Health spending when it goes up by a dollar raises insurance premiums by a dollar, and then we see complete pass through. In fact, higher than complete pass through for reasons we'll talk about into the labor market. You see that a 1% increase in healthcare spending lowers firms' payroll and the count of workers by about ~~fourths~~ tenths of a percent.

And the reason that's such a big effect is because most of us have dependents. So the insurance plan you're purchasing isn't just for you. It's potentially you, your partner and your children. We see that because individuals are losing their jobs. Tax revenue collected by the federal government goes down and then we actually see pretty substantial health consequences. It turns out that losing your job can be devastating for your health.

So let's sort of think about a merger that softens competition. It allows the merging parties to raise their prices. In some ways, we think that the relative

elasticity demand at a given hospital is actually pretty low. And so the deadweight loss in sort of an old school Chicago sense is actually pretty small because quantities don't fall. What this merger really represents is a transfer from consumers to producers.

The vast majority of folks get their health insurance or get their healthcare funded by insurance. They're not paying for it themselves. And so we can think of this increase in price is leading to an increase in premiums of ϕ because this is linked to employment. We can think of this increase in ϕ as raising the cost to firms of retaining a worker by ϕ as well. Now what's critical is when there is an increase in price, it's raising the cost of retaining all workers with ESI, not just those who consume healthcare.

So this worker who's in a factory, who didn't consume healthcare last year is going to become more expensive to retain simply by virtue of having employer sponsored health insurance. We can think that this leads to a downward shift in demand for labor. And then the question is whether this shows up on the intent of the extensive margin, and that's really going to be the function of the sort usual way we think about tax incidence as a function of demand and supply elasticities. And there are all sorts of reasons we could think this is going to be employment, and all sorts of reasons we think the incidence is going to fall on lower-income workers.

So the first is again, this idea of a head tax. It's a much larger proportional change in the cost of retaining workers for lower skilled, potentially lower paid workers. Second, we think this could apply to elastically demanded workers. If there's a range of demand for different worker types, and all the workers become equally more costly to employ, we might think that those lower more elastically demanded workers going to be the ones letting go, let go.

The third is really elastically supplied workers. And here what I'm thinking about is sort of downward wage rigidities. At the extreme end, I'm thinking about something like minimum wage. In practice, I'm literally thinking about my own lab. I don't have any pay flexibility for my research assistants. It turns out the biggest price increase we see post

We're then going to map health spending onto firms. And it's a little tricky because you can't merge the claims data together with the tax data. So what we're going to do is we're going to proxy for health spending by thinking about where each firm's workers live. And we're going to think of a firm's health spending as a product of where their employees get care, which providers, how much care they receive, which we're going to be able to measure in the claims data, and then the prices, the providers where they seek care. So Yale's health spending is a function of how many people from the specific counties where we have employees, which providers they go to care, how much care they get, quantity, times the price of that care.

Then what we're going to have to do is we're going to have to map the price effects of mergers onto that measure. We can't just regress firm's health

The thing that we're going to show you and I don't have time to really build in today is that we can drop huge chunks of our sample. We can throw out the firms that are growing a lot or not growing at all. Firms with high growth, low-wage growth, and we continuously see the same results.

So let's go through our empirics or our data, our results. The first, and this is a log-level regression, so you've got to exponentiate, but a dollar increase in healthcare prices raises healthcare spending by a dollar. And you can think for a fully-insured firm, or excuse me, a self-insured firm, this is mechanical. So this shouldn't be a huge surprise, but it's reassuring that we see it in the data.

Next, we're going to look at health insurance premiums. And again, we have to look at about 5,000 firms because we have a pretty limited data set on that, a pretty limited sample of firms with insurance premiums data. And what we're going to see is roughly dollar-for-dollar pass through. Because we've had to cut down our sample so much to do this, we're going to lose a little power, but we are going to see roughly dollar-for-dollar pass through.

Now what we can measure really well in the data is whether a firm's employees have a health savings account. So you might think with exposure to the higher prices for mergers, firms might shift their workers into health savings accounts. We see no evidence from the IRS data that workers at given firms that are more exposed to mergers are taking out health savings accounts in the tax data.

Now, one of the things we can do is instead of measuring hospital price increases from mergers using these diff and diff estimates, we can actually use willingness to pay. And one of the things that's reassuring is when we rebuild our instrument using willingness to pay, we get almost identical results. Okay, so what about what's happening on the labor market side? What we're seeing is actually slightly greater than dollar for dollar pass through. These are log points, so you've got to shift the decimal place just a little bit. So 1% increase in healthcare prices is lowering payroll and the counted workers at firms by about three tenths of a percent. It's robust to tossing out huge chunks of our sample using willingness to pay instead of our post-merger diff and diff estimates.

And what we see is that all of this is driven by changes at health care firms. And some of the work we're doing going forward is actually looking at where the rents go, what is happening actually at firms in the healthcare sector, and tune in about a year for those results. If you look at our events studies, you can see that these employment changes and these income changes or payroll changes are happening immediately after firms exposure to the price increases from these local transactions. And what's I think really, really important to note here is that these point estimates are equally scaled. It suggests that the changes we're observing are happening on the extensive margin. When firms are exposed to these price increases, they're letting workers go.

Now these point estimates seem really big. And so one of the things we've done a lot in this paper is try to scale our point estimates to other literatures to make

sure they're sensible. And one of the things where there really is a fairly well developed literature is the payroll tax literature. And so we can scale our point estimates to be analogous to a one percentage point increase in payroll. And what we know from the payroll tax literature is a one percentage point increase in payroll taxes. If we look at some of the U.S. studies, the Johnson and all and the Gao and all, we see that a one percentage point increase in payroll is going to get you somewhere between 1.5% and 2% decreases in employment. And that's exactly what we find. And so we're fairly comforted by that. The effects we see scale up with other sorts of payroll taxes, other things that induce costs on firms for retaining workers.

Now, one of the questions we had is are we actually seeing workers separated from the labor market or are we simply seeing the reallocation of workers across firms? And so we're going to take our exposure measure, we're going to aggregate that to the county level and think how exposed individual counties were to rising healthcare costs from mergers in their areas. And so what we're going to do is come up with measures of county level economic outcomes, income per capita, the share of workers who previously had earnings but who had zero earnings, therefore became fully separated from the labor market, or workers who filed for unemployment insurance, again, became fully separated from the labor market.

And what we see is a point estimate on employment that's about a third of the size. So a 1% increase in healthcare spending is going to get you about a little less than a 10th of a point increase in unemployment, which given average unemployment rates or UI rates locally, is actually pretty big. And what it implies is about two thirds of the folks who lose their job at a firm when health spending goes up, find employment at other establishments. One third, lose their job and become wholly separated from the labor market. What we can do is say, who are these workers? Because we can see their incomes in the past and we can bin them into \$10,000 income bins based on their historical income. And what we see is in some sense, reassuringly, we don't see huge changes in employment for workers earning less than \$20,000 a year. Folks who we think are unlikely to have employer sponsored health insurance. But what we do see is effects concentrated among workers earning between 20,000 and a hundred thousand dollars a year and no effects on workers earning over a hundred thousand dollars a year. So it's lower and middle income workers who are the ones becoming separated from the labor market.

And I think this is a sort of critical point. What's the impact of a dollar increase in prices on labor market output locally? And what we see is that a dollar increase

enormous. One year mortality for individuals who lose their job increases by about 50%. And so if we look at the literature over time, we see somewhere in the order of one in 600 to one in 300 of the individuals who lose their job die within a year on average of a traffic accident, ~~self~~ or overdose.

And so the question is, do we see this? We actually see a pretty sizable jump in folks separated from the labor market. Are we seeing them lose their life down the line? And so we're going to bring in CDC's restricted mortality data. We're going to use Case and Deaton's measure of deaths of despair. We're going to focus on suicides and overdoses. We're actually going to leave out alcohol related conditions because we think that takes some time to accumulate. We're going to have a placebo outcome. So in our main outcome, we're going to focus on individuals between 25 and 64. The folks we think have ESI. We're then going to measure deaths of despair among workers over 65 who we think aren't going to be the ones losing their job. We're going to focus on all mortality exclusive of deaths of despair, and then we're going to focus on cancer mortality.

And what we're going to see is that a 10% increase in healthcare spending or a 1% say increase in healthcare prices, it's going to lead to about one additional death per hundred thousand in the area. We're not going to see any changes. There's significant among folks over 65, no overall significant changes in mortality, no overall significant changes in cancer mortality. And once again, we can scale our estimates to give us some sense about whether we should be comfortable with what we're seeing.

And what we see is about one in 173 of the folks in our sample who lose their job die within a year. And in fact, it's two years after the mergers. So it's a year after the job losses which happen after the transactions. So we're a little higher than the literature, and we think that happens for two reasons.

First, we are measuring this at the peak of the opioid epidemic. And if you look at the relationship between job losses and mortality, it's increasing over time as fentanyl and Oxycontin become more ubiquitous. We're right there at the peak, which is why we think this is so high.

Second, the rest of the literature is measuring the effects of job losses. We're measuring the effects of total labor market separations, and about ~~two~~ ^{thirds} of the folks who lose their job, don't become fully separated from the labor market. A third of them do. So if you scale up these point estimates, we're right in line with the rest of the literature.

Now, here's where I think the results are so important for the work that the agency's doing. One of the things we can do with our empirical approach for whole classes of transactions or actually individual transactions is say what the cumulative effect of those mergers are on labor market outcomes and mortality. So we can say, look, the average merger in our sample raises prices by 1.2%. Taking a step back to that first paper, we can say that raises health spending on the order of about \$250 million a year ~~so~~ the 50 mergers that

happen, these one-year effects are about \$250 million. Note that that's bigger than the entirety of the enforcement budget for the FTC, and that those effects are going to persist over time.

Or we can look at the average effect of a merger that raised the HHI by 200 points to a net increase in net HHI over 2,500 points. And what we can show is that those transactions are going to lead to \$16 million in reduced income dropped to zero in the next one, a hundred job losses and about one death from an opioid overdose. And critically what we're going to see is because folks are losing their jobs, because income per capita is going down, we see a concurrent reduction in federal income tax revenue of about \$7 million. And so you can think, look, if this enforcement action, if an enforcement action on average costs about \$5 million, simply measured relative to the gains in revenue to the Treasury from more taxes, more tax revenue because folks are employed, those enforcement actions end up being cost-effective for the federal government.

So what's the sort of summary that because of this mechanical link between what happens in healthcare markets and labor markets, the mergers that we see happening in the healthcare sector are having downstream consequences for workers. Who pays for these price increases? It's workers, lower income it's workers, it's middle-income workers. It's not workers at the upper end of the income distribution. How are they paying for it? They're paying for it with their employment. A small chunk of them are paying for it with their lives. There is deadweight loss from hostile mergers. It's just showing up in adjacent labor market. And it turns out merger enforcement actions are likely in this sector going to be cost-effective. So with that, turn it back over and yeah, we can take some questions. Hi. Hey Ben.

Benjamin Vatter: You asked me to ask a question. So here's my question. Just trying to update together the econometrics of this. And so you're using the merger effect as an instrument for the in prices. Now, when we do, the way that we do a lot of our merger effects is we're trying to really test for an effect. And so we're very conservative in the way that we measure things. We pick control groups, we pick strategies that are fairly conservative, but now suppose it because we do this, we're vastly underestimating the true kind of price effect of a merger. How would that pass through to the estimates that you're showing? Suppose we're estimating, I think if I remember correctly, it was like 1.5% or something like that within this line price effect. Suppose it's like in reality the truth is 5, 6, 7. How does that map through to the assessments?

Zach Cooper: So one, I think a reason to probably take our diff and diff estimates pretty seriously. And then one answer to your question. The first is we actually get super similar estimates when we use willingness to pay. And we think those don't have measurement error to them. It may have some modeling error, but they are so strongly correlated that I actually am pretty confident in our diff

effect of rising spending on employment. But I think given the willingness to pay measures, I don't think that's what's happening necessarily.

Benjamin Vatter:

Zach Cooper: Yeah, it's a great question. I think our work, and we have it in that first paper, doesn't show any evidence of quality changes at merging parties, positive or negative. I think that's consistent with the work Leemore Dafney and others have done. I think the best evidence on the ~~post~~ merger effects on quality is that after hostiles merge, the effects are neutral or negative. And I think our work supports that. Largely what we're seeing is evidence of price increases in the mergers that lessen competition and no concurrent changes, neither quantity or quality.

Speaker 12: All right, thank you.

Zach Cooper: Cool, thanks everybody.

Speaker 12: And that concludes our conference. I want to thank all of our presenters, and in addition to that, I want to thank all of our audience members. None of this happens without all of you engaging with us, asking your questions. And I also want to thank folks at home who may be tuning into the webcast. Thank you for your attention.

Now, if you remember at the very beginning I said there was going to be a pop quiz. So what do we do with our name tags? We're going to give them back to the table where you picked them up from so that we can reuse them again. And for our visitors, what do you do with the lanyards with the plastic FTC visitors badge? Those go to the security desk. All right, two different places. Name tags to the desk where you picked up your name tags and the security badge goes back to the security desk.

All right, and with that, thank you so much. Safe travels and we look forward to seeing you all again next year at our, oh, what number is it now? 18th. It'll be the 18th Annual FTC Macroeconomics Conference.

PART 5 OF 5 ENDS [02:39:08]