



Seventeenth Annual Microeconomics Conference | FTC | November 15, 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, specializing in the empirical analysis of markets in equilibrium. He's the winner of the Frisch Medal of the Econometric Society, was elected to be a member of the American Academy of Arts and Sciences in 2014, and was named the 2017 Distinguished 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 well as environmental and trade policy. Dr. Berry is also affiliated with the Kohl's Foundation at Yale, and is a research fellow at the National Bureau of Economic Research.

Steve Berry: Thanks for that. That was very nice. Sponsoring this conference is one of the really great things I think the Tobin Center does. I think it's very much in the spirit of the kind of work that we want to do. But of course, we sponsor the food. Everybody else does all the other work, including the scientific committee, and the local organizers, and the people who get the AV stuff together and run the day. They're the people who really deserve the thanks, and it's a great job every single year, year after year. I really appreciate that.

In some sense, if they hadn't given me 10 minutes, I would just turn it over to the research, which is the important thing. But let me say a few things about the Tobin Center, and use them a little bit to come back, I think, to the question that 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 pre-doc 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 just talking about healthcare costs, and this is what Zack's going to work on. The healthcare system is a mess. Everybody knows we spend too much money on it. Staffers on both the Democratic and Republican side would actually like to hear something where they would save \$300 million and it wouldn't hurt anybody. What would we do to make that happen?

That's actually an issue where you can get into the weeds. Or if they've already decided to do something, can you try to direct them in the direction of more sensible or less sensible? We have this immensely complicated Inflation

Reduction Act, which is actually a climate bill. The poor IRS is supposed to come up with tons of rulings. One of the things I've talked to them about is, the so-called sustainable aviation fuels made from biofuels, do they actually help the 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 federal government, I've said this before, where the economic staff is kind of unreasonably good.

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 [inaudible 00:10:08]. 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 pre-doc 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 just talking about healthcare costs, and this is what Zack's going to work on. The healthcare system is a mess. Everybody knows we spend too much money on it. Staffers on both the Democratic and Republican side would actually like to hear something where they would save \$300 million and it wouldn't hurt anybody. What would we do to make that happen?

That's actually an issue where you can get into the weeds. Or if they've already decided to do something, can you try to direct them in the direction of more sensible or less sensible? We have this immensely complicated Inflation Reduction Act, which is actually a climate bill. The poor IRS is supposed to come up with tons of rulings. One of the things I've talked to them about is, the so-called sustainable aviation fuels made from biofuels, do they actually help the

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 federal government, I've said this before, where the economic staff is kind of unreasonably good.

You guys are better than we deserve. Really, you're fantastic. I'm not sure we deserve that, but it's because there are people who want to do good. There are people who will turn down an outside opportunity in order to do work that makes the world a better place, and that's super inspiring. We've gotten people who have been in the White House, who have been legislative policy directors for the White House, things like that, who want to come and work for us because they see that we're sincere and they actually want to get that done.

I think it's really important for us to maintain this connection between academic economists and the government economists. I think we have complementary abilities in this regard. One is if we have an idea, we have to figure out how are we going to get a seat at the table? The government staff are lucky. You guys have a permanent seat in the building, and that's something that you're really lucky about. And you know the policy questions, because they're coming at you fast, so you don't have to wonder about what's the important policy question? It's coming right at you. I think those are things that are valuable for us, that if we have people who are in the inside, who know who we can talk to, that's absolutely fantastic.

Now, I think one of the advantages that we have is that we can pivot. If antitrust is not the issue of the day, even though we might have some really good antitrust work going on, we can do pre-K policy. I think what's interesting about this conference is that a lot of the research that people here do, a lot of the research that the government economists present, is relevant to maybe not the question they're being asked today. Maybe they were asked that a few years

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 non-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 pre-K on 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 because of how it's done. I think we have to be careful about that. In an age where experts are distrusted, we worry that a bunch of different political movements have said economists are worthless. I think keeping that credibility is something we can do. And actually, again, I'm sincere, and corny and stuff. I actually think it'll pay off in the long run. Why don't we hear some actual research?

Thomas Koch: Good morning. I'm Thomas Koch. I'm an economist here at the FTC. I have the privilege of introducing the papers selected by Zach Cooper. The first one will be presented by my colleague, Miriam Larson-Koester.

Miriam Larson-K...: Sorry, how does this work. Oh, the big green button. There we go. Great. Yeah.

Thank you so much for having us here today. It's truly an honor to get to present this work at this conference. This is joint work with my co-authors, Dan Hosken and Charles Taragin, who are sitting right here. The usual disclaimer applies here: this does not reflect the views of the Federal Trade Commission or the Federal Reserve board. We also don't have any financial disclosures on this paper.

With that, I'll jump straight into the research question. We're going to look at downstream horizontal mergers. And the question is, how and when do they affect consumers and workers? The idea is that mergers that may be good for consumers, could be bad for workers, or good for workers and bad for

consumers. And we want to figure out in what case that occurs, and what the interaction is between the consumer market and the labor market.

To do that, we're going to take a two-level vertical supply chain model. This was originally from Horn and Wolinsky, more recently from Sheu and Taragin. The way it works is upstream, there's collective bargaining over wages, so efficient bargaining. And downstream, there's a differentiated Bertrand. This is going to allow us to simulate consumer welfare and worker welfare at the same time. There are some key mechanisms in the model. The first one is going to be just think of your standard downstream merger simulation. You have product-market recapture. That's going to tend to increase prices and decrease quantity. That's going to affect workers, because decreasing quantity means decreasing employment. But in addition to that, we're going to layer on this bargaining model. The first additional mechanism that we have is actually that as firms become more profitable due to the downstream merger, there's an increase in the bargaining surplus, and this can potentially increase workers' wages once we have that bargaining game added on.

The next effect that we're adding in is going to be an increase in the firm's bargaining leverage. You can think of this as if a single product firm is negotiating with its workers, if it can't reach an agreement, it's not going to be able to produce. But a multi-product firm will be able to recapture some sales, through downstream consumer substitution, into its other products. And therefore, its threat point in the negotiation is going to be higher, giving it more leverage and potentially decreasing worker wages.

The fourth and final mechanism that we have, and maybe the most important, is the change in worker bargaining leverage. To the extent that there's a change in the labor market concentration, we're going to adopt a sort of static version of the Jarosch, Nimczik, Sorkin paper and say the employers will not bargain against themselves. The idea is that if I'm a worker, and pre-merger, I can negotiate with firm one and firm two, I'll receive offers from both firm one and firm two and use them as threat points against each other. Post-merger, if firm one and firm two merge, I'm no longer receiving two separate offers, so I lose one of those options in my outside option. This will lower my threat point, lower my leverage, and potentially decrease my wage.

Another key aspect of the model is that it's pretty flexible in terms of thinking about whether labor markets and product markets are the same. Especially geographically, we're thinking about potentially labor markets could be broader or narrower than the product market. The first case is thinking about a service industry. It could be the case that consumers are willing to not really move very far to consume products, but workers might be able to commute a lot farther. For example, you could have two hospitals on other sides of town. In terms of consuming especially acute care hospital services, I might just go to my local hospital. But a nurse might be willing to work anywhere in the commuting zone, and go across town to work at another hospital, so the labor market could be broader. Secondly, you might think about in a tradeable goods industry, the

product market could actually be a lot broader than the labor market. You could think about maybe automobile manufacturing. There could be plants in different states. They're both producing the same car into a national market, but those local geographic labor markets are not interrelated at all. Implicit in that is that mergers can occur within or across geographies. In particular in this paper, we're going to focus on cross-geographic mergers a lot, because we're thinking about completely shutting down the overlap in either the labor or the product market, to focus on how the effects on workers can differ, so you can think about workers only being affected through the product market, or consumers only affected through the labor market. That's an important point, so I think I'm going to go into a little bit more detail on exactly the types of configurations we're looking at in order to illustrate the model.

Our baseline configuration is what you might think of, is where there's perfect overlap between the labor and product market. You can think of this as a rural area where, say there are two hospitals. Everyone is both working at those two hospitals and consuming at those two hospitals, there's not a lot of other options around. And so in this case, the change in concentration is going to be the same in the product market and the labor market, and all of those model mechanisms are going to be in effect.

The next case that we look at is a case of a broader labor market. In particular to illustrate the model, you can think of there being more or less overlap in the labor market versus the product market, but we're actually going to take the case where the merger is occurring across the geographies in the product market, so there's actually no change in product market concentration from the merger. Thinking about the hospital across town merges with the other hospital across town. That doesn't affect the demand for hospital services, but it does affect the workers in the commuting zone for those hospitals. And we do this to illustrate the mechanisms of the model, because in this case, the only impact on workers will be through the labor market.

The third and final case we take is the opposite of that. Here we have the tradable goods case, so we're going to look at mergers across geography, like manufacturing plants, selling into the same downstream tradable goods market. This is, again, an extreme case, taking the cross-geographic merger where the only impact on workers is going to be through the product market, and that downstream product market overlap. In terms of our approach, we're going to take a simulation approach. This is partly because the model is extremely flexible. There are a lot of parameters, a lot of configurations to cover, and so we want to try to calibrate it to some real-world industries to try to discipline what results we're going to get, in terms of seeing predictions that are actually relevant to potentially cases that we would see. We're going to take some observed shares, wages, margins, and costs from real-world industries and run simulations of a broad set of fake mergers to get a sense of how the model works.

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 only on product market does actually pretty well at predicting when mergers will harm workers, as long as there is product market overlap. But in that case where there's only labor market overlap, in that case, obviously the downstream simulation cannot do a good job.

In terms of where we fit in the literature, obviously there's been a lot of recent empirical evidence on labor markets not being perfectly competitive. And recently Prager and Schmitt, and David Arnold, the impact of mergers on wages. But some of the more classical models of monopsony don't take into account strategic interaction. Obviously, the Jarosch paper we saw earlier is an exception, and also the Berger et al paper that we saw at this conference last year. The Berger et al paper is pretty close to ours. They feature local labor market competition and perfectly competitive product markets. We're going to be complementary, in that we have both competition in the labor market and in the product market. And we're building more off of the merger simulation literature, so differentiated Bertrand, Werden and Froeb, and Nash Bargaining with GNT.

With that, I'm going to jump into the model, and talk about first how the bargaining works and how wages are set, since that's the key focus of the model. We have here, I think we've already seen this maybe a few times at this conference, but a standard Nash Bargaining function. We're choosing a wage that's going to maximize the product of the gains from trade of the worker and the gains from trade of the firm, weighted by the bargaining power parameter, λ . As λ goes to one, the outcome for the firm will improve. I'm going to now go into detail on how we specify each of these payoffs. That's sort of how we're going to get at the model mechanisms.

The firm's agreement payoff is going to be a pretty standard profit function. The key thing here is that we're assuming a Leontief production technology, so one worker is producing one unit of output and using one unit of cost as well. That ratio, one-to-one, is not necessarily have to be that, but we choose that for simplicity. One consequence of that is that we have constant marginal costs, which is a pretty standard assumption in merger simulation.

The firm's disagreement payoff is going to be zero if they can't produce, if don't reach an agreement with the worker. But we have denoted here as Z_I , the set of

other products owned by that firm. And so that's going to be the consumer substitution, the recapture, that allows the multi-product 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 post-merger, the firm's threat point is going to go up.

I won't talk too much about the downstream, since it's a standard Logit demand system. I will say that the margins of the firm also depend on that set, ZJ, and so that's how we connect to that second mechanism in the model. Post-merger, the larger the set ZJ is, the higher the margins are, the more profit, that will tend to increase wages.

For workers, we have a labor market size that may be different from the product market size, so we do allow for different market elasticities for workers versus consumers. But there is this Leontief production technology, and so you need to write both worker shares and product market shares in terms of quantity, which is by scaling them by the market size, which is given by that equality there. The worker agreement payoff, we assume that they get the wage bill within the market, so just the wage times the quantity. And that, again, can be written either in terms of worker shares or product market shares.

The worker disagreement payoff is going to be the key thing. We assume that if the worker doesn't reach an agreement with the firm, they can work at other options within the market or outside of the market. If you're thinking about nurses in hospitals, this would be the nurses can work at other hospitals in the market, or they can go to work at a long-term care facility or doctor's office, which would be considered outside the market. Again, the key thing here, again, is that set ZJ, the set of products owned by the merging firms. We're assuming that if two hospitals merge, I no longer get an offer from both of them. I'm only going to get an offer from one, the merged firm. And therefore, I'm going to have fewer options inside that disagreement payoff, and more people will be diverting to the-

PART 1 OF 5 ENDS [00:32:04]

Miriam Larson-K...: ... Is inside that disagreement payoff and more people will be diverting to the outside option.

So now I'm just putting back together all of those components into the larger bargaining game. So we have here again that the worker payoff is the wage bill minus the share weighted diversion to other options in the market. And the firm gains from trade is their profit minus their recapture in the downstream product market. And so the wages are going to solve this function as weighted by the bargaining power parameter.

In terms of solving this, we make the standard Nash-in-Nash assumption. So when each negotiation is happening, they're holding fixed all of the other wage

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-the-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 take the skilled workers wages, again, because we think skilled workers might be having more specific human capital. And then we're going to look at local geographic labor markets across Colombia and assume a national product market. We do our best to exclude imports and exports and unfortunately we have to assume that all the plants are independent pre-merger because there's no data on ownership in the census.

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 across geographies. So no change in labor market HHI, but focusing on the effect of the HHI in the product market, which is the broader tradable goods market. We also limit to meaningful mergers. So we're going to only look at mergers within HHI over 1500 and Delta HHI over 100.

So finally getting into our results. For the first configuration, we have a set of 855 mergers. We're going to focus in the interest of time just on the changes in worker surplus from the mergers. These are box and whisker plots that cover that distribution of simulations. So the median is going to be that line in the box. The edges of the box are the 25th and 75th percentiles of the distribution, and the tails are the fifth and 95th.

In the downstream-only model, again in this model, the only impact on workers is through their employment. We do see a pretty negative impact on workers. But moving to the middle model, again, this is adding in both the bargaining surplus effect and the firm leverage effect. So the predicted impact is not clear from that because the bargain surplus will tend to increase wages and the firm leverage will tend to decrease wages. And so we actually see that workers are kind of the same or maybe even a little better off moving into the bargaining model where their wage can adjust up or down.

But then when we move to the full model and we add in that worker leverage effect, we see a pretty stark decrease at least in the median of the worker's outcome. And so we can conclude from that that the worker leverage effect is

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 overlap. The only change in the workers' threat point is going to be through equilibrium effects. So the wages at the other hospitals in the market will change as a result of the merger. And you see this has an ambiguous effect on workers. The median goes down slightly, but the variance goes up a lot just because the wages at their other options are changing in the market.

The next set of results I'm going to show are the relationship to Delta HHI. So we want to get a sense of how predictive the Delta HHI is of worker harm. There are a couple of other reasons that we want to do this, and one is that the distribution of mergers differs a lot by the configuration. So in order to compare apples-to-apples, we want to look at the distribution of effects by the Delta HHI. So you're looking only at bad mergers or good mergers.

The next reason is that the model mechanisms might vary with Delta HHI. You may see some benefits at lower that disappear as you go to higher, et cetera. In order to do this, we're going to take some common measures of merger impacts, so the more intuitive measures like wages, prices and output, and see how they vary with Delta HHI.

Here's the first set of results for that. So each in this graph is a merger simulation, and on the X-axis is the Delta HHI. This is for a first configuration. So here the Delta HHI is both in the labor market and the product market. It's going to be identical for both markets, so it represents both. The thing that I want to point out here is that there seems to be a very strong correlation between Delta HHI and both wages and output. And since both wages and output are affecting the workers surplus, the workers surplus is also very strongly negatively correlated with the Delta HHI. And in fact, the price tends to be a bit more dispersed in terms of its correlation.

Moving on to our case where there's only labor overlap and no product overlap, we can see a very, very strong correlation between the wage and the Delta HHI and not as much of an effect on price or output. In a sense, this makes sense because there's no product market overlap here, so you wouldn't expect a first-order effect on price or output from this merger.

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 pass-through 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 pass-through 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 pass-through.

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 pass-through 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. And finally we looked at the conventional product market screening tools and found that they're actually pretty effective. And 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.

And 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 right button on here? Thank you. Okay, great. So when Tom invited me to discuss this paper... Thank you by the way, and thank you to Viola and Sam for getting us all organized. When I was invited to discuss this paper, I expected to have to defend to an audience full of antitrust economists why you should care about how we model labor markets.

After yesterday, I'm not sure that's really necessary, but I will point out for the uninitiated that the new merger guidelines that came out at the end of last year for the first time explicitly call out labor market competition as a component of merger review. And there's been a lot of discussion about this. There's a serious question that Miriam raised quite clearly toward the end of her talk about whether, in fact, just using the standard screens that we already have on the product market side would already catch most of the transactions that we might think might be harmful to workers.

And if that's the case, then smarter people than me have pointed out, Nancy is in the room over here, that enforcement is not costless. And if you're going to add a labor screen to every single transaction that you look at, you may end up being a less active enforcer overall because unfortunately we don't have the resources to just double the staff at the FTC or the DOJ as much as we might like to.

So this paper helps us to really think about the stuff seriously for the first time. It's not an understatement to say that putting together strategic interactions on the product market side and strategic interactions on the labor market side in

the same paper is not an easy feat. I have thought about this and come up dry, so to Miriam and her co-authors, thank you for moving us in the right direction.

They have product market competition that I think all of us recognize as being pretty natural, differentiated products, Nash-Bertrand. That's layered on top of a Leontief production function. I'll come back to what the implications of that are in a couple of minutes. And then on the labor market side, we have negotiations between an employer and something that looks vaguely like a union. It's this group of workers and they're negotiating also Nash.

So the key results of the paper pulling out just the versions of the simulations that use the full model are twofold, or rather there are many results, but there's two in particular that I want to highlight. First if you look at the panel on the right, these are the simulations from the Colombian Manufacturing data where there is no labor market overlap or at least that's what the exercise is meant to represent. You can still get negative effects on the labor market even when the overlap is entirely on the product side.

The other thing that I want to point out here is that you can have worker effects that are larger when you have both labor and product market overlap than when you just have labor market overlap. Those are the two panels on the left. If the magnitudes I just said don't make sense, just note that the Y-axes are different across the version that has both labor and product market overlap versus has just the labor market overlap.

So what this paper tells us is that the product-based screens do a pretty good job for the kinds of transactions that we are already used to thinking about. In other words, if you are thinking only about traditional transactions that primarily impact product markets, then seems like we do a pretty good job already of also catching harms in input markets.

The question is then, well, what fraction of transactions is that? Right? And the paper tries to make inroads in that direction by saying, well, let's look at what would happen if you actually looked at transactions that primarily are labor-impacting rather than product-impacting. And there they find, kind of mechanically but kind of not, that the traditional screens, of course, we'll catch none of this.

It's really hard for me to overstate just how timely and important this is. This is a serious discussion that obviously the agencies are having. It's a discussion that lawyers and economists are having from outside of government and it's really important for us to figure out how we should be targeting these enforcement resources that are limited. And so the paper is a really nice step forward in helping us to understand those things.

It also still leaves a question open for the rest of us of okay, we've now established, thank you Miriam, that there is a class of mergers for which the

existing screens don't do a good job. How many transactions are in that class? And when I conclude, maybe I'll have time to come back to a call to action for the rest of us. In some ways, enforcers have gotten in their own way for measuring this because the most recent update to the HSR pre-notification form could have included a labor screen and didn't, which means that we're now not going to collect systematic data about how much labor market overlap there actually is among proposed transactions, which I think is kind of a bummer that we won't get those data. And as a result, maybe the rest of us in the research community who are not bound by agency rules can step in.

Okay, so I'm going to do the reverse of what Heski did yesterday because I'm a pretty poor theorist. I will focus most of my comments on the simulation exercise, but I do have one comment on the model, which is that we're using a Leontief production function here. That's important for tractability and we want to maintain tractability. It does have some implications I think for how we interpret the results of the simulations, right? Because in particular what you're getting here is a disagreement payoff for the employer that's pretty severe. They don't get to produce that particular product at all, even if there's some recapture through a multi-product firm, which means that you basically need something else in the model, some other parameter, to rationalize observed outcomes that look pretty good for the firm.

And it's possible that that's what's driving the estimated parameter for the firm bargaining weight being so high, which is something Miriam didn't really have time to cover in her talk, but in many cases it seems to be pretty close to one. That, in turn, might explain why you don't get very much pass-through to workers, which is something of course that you did mention. So I think we want to be a little bit careful in what we try to extrapolate from the simulations, bearing in mind this Leontief assumption.

Let's also talk a little bit about the worker-side results. So one of the figures that I put up showed you this comparison between labor market effects in markets... Or sorry, in transactions that involve both product and labor market overlap versus ones that involve only labor market overlap. And I was surprised to see smaller effects on workers in the latter case because you would think that there would be some offsetting effect when there's also an increase in product market power, you get increased prices. Some of that ought to be passed through to workers.

Maybe it isn't just because of the high firm bargaining weight. But another possible explanation for what's driving this in the calibrations is that you have actually different geographic definitions of what counts as the labor market across these two sets of simulations. The second set uses those broader labor markets that Miriam mentioned. The HRRs. On average, HRRs are something like 10% larger, or sorry, 10 times larger than HSAs, which means that you're not fully doing an apples-to-apples comparison.

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 map of HSAs. This is the DC area. Many of you probably work here. I imagine many of you live outside the district and yet Arlington and Bethesda and the district are all separate HSAs.

So I think we maybe want to be pushing toward HSAs that are a little bigger. The analysis already does this. I should be careful to say they don't try to do anything in HSAs that are super tiny, but I think maybe subsetting some more or maybe expanding the product market definition to something that looks more like an HRR might be warranted.

So with that apples-to-oranges comparison, I think it becomes a little bit hard specifically to do just the exercise of do we think that workers are going to be more affected when there's product and labor overlap as opposed to just labor overlap that doesn't invalidate either of those exercises in isolation, right? It's just an issue of whether we can compare across them.

Clearly there are reasons to be concerned about worker harms in both cases as the simulations show. Trying to find a setting where it's very clear that you just have labor overlap without having product overlap is really difficult, which is why they've had to set it up this way. And I'm very sympathetic to it being hard to find not just an industry, but also an industry with publicly available data, that would allow you to do this even more convincingly.

Honestly, I wonder if the answer is just to simulate an industry. If you're already calibrating a lot of things and simulating a lot of things, maybe just come up with one that fits the needed structure.

Okay. I will more or less stop there. I just want to emphasize again how incredibly timely this is. As a discussant, it's my job to quibble with some of the implementation choices, but this is a hugely important policy concern right now. We don't know a lot about it. And really I think it's our job as a research community as well as an enforcement community to try to figure out what fraction of transactions are actually just completely flying under the radar but may be harmful to labor markets because we haven't been implementing those screens.

I'll stop there.

Speaker 2: And now we have time for some questions.

Miriam Larson-K...: Thank you so much for the discussion by the way.

Audience Questi...: Hi. I enjoyed the clarity of that paper because there's a lot of things going on and having something that actually lets you work through whether the different pieces without it being kind of a gigantic model was really helpful.

So the one comment I'd have is I've always thought the difference between a model with a union versus a model with workers Nash bargaining was are we having negotiations about the marginal product of a worker versus the average product?

And it just strikes me that this would be an easy switch to go from one type of labor market to another depending on what's being split. And given the prevalence of unions in the United States being small, it might be simple to address those two kinds of setups in the paper.

Miriam Larson-K...: Yeah, that's a really good question. We've actually thought about that quite a bit. We have this constant marginal cost assumption, and so I think we have to go through the math for sure. But I think it's actually the equivalent problem if you had individual workers bargaining in this case because the marginal worker is the same as the entire worker in terms of the math, so we'd have to put in some sort of diminishing returns to scale or convex costs, I think, into the model.

And then we could look at thinking about the marginal worker being maybe not as pivotal to the firm in terms of their output.

Audience Questi...: Hi, I'm over here.

Yeah. So I'm wondering, so how exactly you deal with worker heterogeneity in the sense that you could imagine that more specialized workers are more hit by a merger that I guess is in the same area, but at the same time they might also be more mobile. And I guess for some workers, minimum wages might be binding and I'm wondering whether you'd take that into account.

Miriam Larson-K...: Yeah, we definitely can adjust that worker's outside option, right? So we do set it according to what it is for nurses and what we think it is in Colombia, but it could be much broader for lower skilled workers. In terms of the minimum wage, we have not thought about that, but maybe we can talk after about...

Audience Questi...: So this is kind of on a similar note to Alan's question, but a little unfair I think. But so this stuff on using modern tools on strategic interactions to think about labor market seems potentially really valuable. But there's also this older literature in labor economics taking really seriously how to model labor bargaining. And I fear if we don't read it, we're going to end up reinventing the wheel and having a lot of debates all over again.

So they, for example, think a lot about what are labor unions or labor side exactly negotiating over and what do they care about with it and without an

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 Larson-K...: 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 Larson-K...: 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 vertically integrated hospitals and non-vertically integrated insurers, and, sorry, the complete network of agreements. We have very detailed administrative data on ownership, on insurance plan structure, on networks, on insurance rates, on enrollment and premiums and individual's incomes and how much they pay out of pocket versus how much the insurer pays and so on.

And so the paper is structured into two parts. The first part that I'm not going to talk too much about today because of the sake of time is largely descriptive.

And I tried to ask, are vertically integrated firms different than non-vertically integrated firms? We have three headline findings from there. One is that vertically integrated hospitals charge 13% less to their own insurers than they do to rivals conditional on the patients' demographics and the case and the complexity and the treatment and all the other things that you would like to control for that might affect the cost of the service. And by charge here, I mean the true price. They charge [inaudible 01:06:43] 13% less to their own insurers.

The VI insurers are very much in the business of getting people to stay within their systems, and they do so to a large extent by providing them generous deals in terms of cost sharing. So cost sharing is more generous if you stay within the system. And despite our best efforts, we fail to find any evidence of differential quality, cost or treatment decisions at integrated hospitals for their own patients where the literature on integrated care will tell you there should be something happening. This is where they have some informational advantage and they don't seem to be leveraging it, which is consistent with some other findings within this literature on integrated inpatient care.

What I'm going to focus in today is on our findings from a model that we develop an estimate of the market, which includes endogenous hospital prices and planned premiums and most importantly, endogenous plan designs. So we endogenize the structure of the plans, meaning the networks and the coverage that is offered by each plan. And it turns out that that's really important. So we're going to find that vertical integration reduces welfare by around \$42 million. To put that in context, it translates to around a 6% increase in inpatient spending. Now, inpatient spending is very inelastic, so 6% is a large increase for a market. And to benchmark it even further, you can think that for the people that are going to get harmed by vertical integration, this is equivalent for them to about five monthly premiums of surplus harm.

And the reason why it's going to happen is because vertically integrated insurance networks are going to silo the VI hospitals from their competitors and those competitors are going to face less competition and going to increase their prices. But importantly, this happens only through the endogenous restructuring of insurance networks in response to the market structure. And if you fail to account for that endogenous restructuring, you would get the opposite side on the welfare effects of VI. So that's I think one of the interesting things we find here.

So to highlight the mechanism, I'm going to start with a simple illustration. This is the simplest setup that you can put this into. There are two hospitals upstream, H1 and H2. There's two insurers downstream, MA and MB. And the market has to determine by various means prices, which are denoted by P , coverages, which is one minus the co-insurance rate, this is... You'll see in a bit. It's the only meaningful cost-sharing element in this market, which is denoted by C , and the premiums. And the demand is structured by the arrows, represented by the arrows. And H1 and MA are integrated, hence they're both in blue. And what it means is that they set their own price, $PA1$, and the

coverage of the insurance plan, so CA1 and CA2, the premium to maximize their joint profits, while MB is not integrated with H2 and it needs to determine its own coverage rates, so that's the design of its plans, and its premiums unilaterally to maximize its own profits and it has to negotiate all the other prices.

So I want to highlight the distortions that VI introduces or the changes that VI introduces at the top. And the paper will include also the changes in the demand and the premiums and whatnot, but just the key forces at the top here. So the first thing that vertical integration does is that it creates an internalization of integrated profits. The insurer internalizes the effect of each plan design on retaining profits within the system and it creates these incentives to self-preference, so to increase coverage within the system, and to narrow its networks, which in this setting is represented by lowering the coverage outside to rival hospitals. The hospital, the integrated hospital, understand its effect of its prices in the demand for insurance for its integrated system. So that eliminates or reduces the double marginalization, pushes its own prices down, increase incentives to increase rivals' costs with the attempt of rising the cost to the rival insurer will shake off some of its demand and hopefully some of that will come to my own insurance plan.

But the interesting action happens to the rival insurer. So the rival insurer is faced in a position in which these integrated insurers are going to offer pretty generous cost sharing terms towards their own hospital. So they have pretty generous in terms of risk protection that they're competing against and they need to offer something competitive in this market and they need to decide where are they going to provide this risk protection or this coverage to. Now they cannot do so or they cannot... They're dissuaded from doing that at H1 because H1 has this incentive to raise their costs. And if they provide generous coverage at H1, H1 is going to take it as an opportunity to increase their cost substantially, make their plans look bad and get a lot of their demand. So they're going to have to do that at H2. They're going to have to lean on the non-integrated partner.

And you see that as this is happening, the arrows in the middle and the diagonals get fainter. Less and less demand is flowing within them. And so that is siloing. That's the silo of hospitals that is happening. And in a static picture, this looks a lot like foreclosure, but it's coming from the downstream. It's the voluntary decision, strategic decision, of the downstream to change in a way its inputs, rather than the upstream foreclosing the downstream. And that comes with differences in the analysis and the outcomes. But fundamentally, our finding, the way to read it, is we put this through the empirical lens and quantify things. What we find is that the losses in access from this diagonal lines getting fainter and the access through it becoming more difficult and the increased price on the gray line, the non-integrated line just offsets the gains from lower prices along the blue line, along the integrated line. This is fundamentally what we're finding. So here's a crash course on the Chilean private healthcare market. So we're looking at it in Santiago within 2013, 2016. On the insurance side,

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 plans for different age groups and gender and family status, and there's community pricing. So basically if you're a woman between 18 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 maximum out-of-pockets. Really for inpatient care, the most relevant thing, or fundamentally, the only relevant thing, is this coverage rate, which is one minus the co-insurance, which is a fraction of the bill that the insurer will pay if you go to a hospital. And the regulation dictates that insurance networks have to be tiered. Every general inpatient care hospital has to be in either a base tier or preferential tier. The base tier dictates base coverage. The median is around 55% of the bill is paid by the insurer. And at every plan, every insurer can decide to have a different set of hospital to call them preferential and to have a more generous coverage, so 88% of the bill. And that effectively on the ground works a lot like having in-network and out-of-network providers with the caveat that there's going to be negotiation for what we would call out-of-network providers here, so it's like the prep-based tier hospitals, and that every insurer can offer various structures of this kind of in-network and out-of-network, or preferential and non-preferential.

On the hospital side, there's 11 large inpatient hospitals providers during our time in Santiago. They provide meaningfully heterogeneous quality. And that really matters because our upstream model, our upstream providers here, are heterogeneously differentiated or horizontally differentiated in meaningful ways. So some of them are the best at providing maternity care. Others are good at infectious diseases. And others are good at oncological care. And you as a single patient might be at risk for different types of demand or different types of treatment that are better delivered at different types of hospitals. Importantly for us, there's two-star hospitals that are non-integrated that provide almost uniformly better quality across the board. And you should think about them as those are the ones that are benefiting the most from getting excluded or getting isolated from high-quality integrated hospitals.

There's an outside option here both on the public side for insurance and for 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 each own three hospitals, and there's one set of hospital insurer pair that is integrated at the beginning of our data and then stops being integrated. They get into a fight around surplus splitting, and then they stop being integrated. We leverage it very minimally, and I'll show you

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 non-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 maternity care and others are better at oncological care.

What is important about this model is that we have this differentiation, so there is value for broad access in this model, both because of location differences and because of horizontal differentiation in hospital care. And there's heterogeneous price elasticities, which means that you should expect people to buy different coverage levels. On the enrollment side, people are buying plans for themselves and for their family. They care about the premiums that they have to pay. They care about the surplus that their whole family unit is going to derive from that, and that has to do with the risk that they're exposed to. They might also be affected by other characteristics of the insurance plans. These insurers market themselves slightly different to different age groups, so that shows up in the data. But most importantly, because families live in different places in the city, they're exposed to different types of risk, they value

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 under-provision 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 coverage that they're going to set at the base tier and the preferential tier and who are they going to call preferential at every plan that they're offering. Now they're going to do this taking into account the equilibrium consequences of this decision, so how does it affect the sub game, and some design costs, which largely have to do with slacking some regulatory constraints.

Now, this was the slide that was the hardest for me to cut off from this presentation, which is about how we go to solve this problem, because this is a rather complicated combinatorial optimization problem, but it turned out that it falls within a class of combinatorial optimization problems that the combinatorial optimization people have made a lot of progress on since the 2010s. And so we are able to leverage some results that in fact speak a lot clearly to the mechanism design problem that these insurers are trying to solve and it works remarkably well at solving these problems. And it works at scale for 100s of plans for each insurer. So headline estimates, much like in the US, we find that people are substantially more elastic to premiums than to hospital prices. Unlike the US, we find that hospital price elasticity is not nearly zero. So it's -0.8 for people that are not from healthcare. This is high. In healthcare, it's basically zero. Now, the reason why it's more than zero in Chile is first of all, it's

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 re-optimize their coverage, their plans. We let hospitals and insurers renegotiate their rates, insurers to re-optimize 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

coverage is 52% in the baseline and 78% in the baseline. And this falls on the base by 6% and increase by 2% on the preferential. And it's even same pattern, but even starker for the non-VI insurers.

So what is happening here is that the gradient of moving, staying in network versus going out of network is becoming stronger, and so steering is more effective. They're steering people in a particular direction in a stronger way. And you'll see exactly why. In terms of who gets to be on the preferential tier, in baseline, the way to read this number is self-preferencing states of VI hospital how likely it is to be preferential on a VI insurer plan. It's 67% of the time in baseline, and that drops by around a third once they stop being integrated, and it drops by around a half on non-VI insurers. So the rate reads 22%, 22.1% there on the non-VI insurers. Other VI preferencing is how likely is a VI hospital to be preferential on a non-VI insurance plan. And that drops from 22% to around 10%.

But the most interesting action comes from this 7.4% on the Star hospital preferencing rate on the VI insurers. This seems small, but that's really, really important because that's the margin in which VI insurers are going to pick up preferential plans to the Star hospital, high quality, non-integrated hospitals, and they're going to put it for the people that really value that high quality, high complexity, high value care, and that's where it's going to create competition for their former partners, high quality partners, and going to affect the prices.

So moving on to prices. One thing to note here is that because of all of the steering, the average price in the market and the average price experienced by consumers are very different. So when you eliminate vertical integration, double marginalization kicks in and the average price between H1 and MA or this is meant to represent the average integrated price, that increases, increases by 23%. And the downstream insurer no longer has an incentive to depress its premiums to steer people to its own hospital. So that also increases and also faces higher cost. But the insurer is not passive. It's steering people away from where this double marginalization hurts the most. And in fact, along this vertical chain, the average price experience among consumer falls.

And what is happening is that all of these... If you remember, they own three hospitals, and you should think about them as vertically differentiated in quality in a way for simplicity. They're taking away people from the middle ground. They're steering low-validity, low-complexity care to the downtown hospitals that are cheaper, that are low-complexity. And it's reducing transportation costs. So that's good for consumers. And high-complexity care, high-validity care, which carries the bulk of the surplus, it steers towards the high-complexity, high-quality former partners that now face increased competition from the Star hospitals. So they face lower prices there.

On the other average prices, one thing that is important here is that eliminating vertical integration increases competition both upstream and downstream, but upstream providers are meaningfully differentiated organizations. They provide

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 non-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 low-value care to downtown Santiago for 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 re-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 non-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 non-VI 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 quarter of all their profits, 87% of it is recaptured by their rivals, and non-VI insurance profits increased by around 19%. Now, consumers are almost all better off. So 72% of them are better off. Who is worse off? Well, the ones that are very loyal to the VI insurance because they see double marginalization kick in, but they are losing around \$8 a member on average and the ones that are benefiting are getting about \$52. And they are the ones on non-VI plans. And for them, that translates to around a five monthly premium surplus gain.

But now the important part of this is well, VI is welfare decreasing. This is the net. The welfare number here is negative or positive because we banned VI. But if you held coverage fixed, you would get a different answer. So what I'm showing you here, the orange bars are the full adjustment numbers, which is what I just showed you. The green bars is what would happen if we held plan design as we see it in the data. And so what happens is that if you hold it as you see it in the data, you would find that VI is largely a good thing because what are you introducing? By eliminating VI, you're introducing double marginalization and you're eliminating the rising rivals' cost effect. But the rising

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 non-VI 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, really large. They're like on around 20% cost efficiency gains. So that's not

reasonable. But quality effects could offset this, but they would still have to be quite large.

And we also benchmark this to what would happen if people had different preferences. Preferences might look more like the US preferences. And I could tell you that whenever the elasticity to premiums is larger than to prices, it more or less looks the same. And the antitrust implication or the policy implication of this is very clear. If you want to curtail the negative effects of VI, you need to regulate the networks because that's where it's coming from. And we know how to regulate networks, or at least we have experience with it. And so this is where it's coming from. That's it.

Speaker 3: To discuss the paper, we have Shoshana Vasserman from Stanford.

Shoshana Vasser...: All right. Thanks very much for... Let's see. Hold on. It's this button? Okay. We have to go through all of the appendix slides.

Speaker 4: There's so many.

Shoshana Vasser...: Oh, boy. Still appendix slides. Oh. One second. Are we? Okay. Is there a way to click through faster? You guys are seeing all this stuff? Okay. Oh. There we go. All right. Thanks so much for having me.

There's a lot in this paper, and I knew there would be, as soon as I saw that the starting point was the Ho and Lee paper. So I anticipated, I think, correctly that much of Ben's discussion was going to be about the modeling framework and thinking through the logical steps and how they break it down in the paper. As you saw, they broke it down in a very detailed way at every step of the way. And so I thought that instead of spending a lot of time quibbling, the most valuable way I could spend my 10 minutes is by giving you guys some intuition about what's really driving this underneath. As you could tell, there's lots and lots of data and lots of careful work to understand why the modeling at every step of the way actually made sense. I'm going to give you just a snapshot of some of the descriptives and some of the work that was in the paper to justify why the decisions made sense in fact.

Okay. So Ben gave you an example of the market structure and a dummy example, here's a figure from their paper that shows you exactly the market structure. So there are, I think, nine hospitals here. They're labeled... Or I think nine. I don't know. Eight? Eight hospitals for some reason labeled H2 to H9, or H10, H11. Never mind. There's eight hospitals. These are all private hospitals. And there is one public hospital. And there's several insurers. And the dashed boxes are around sets of hospitals and insurers here are going to basically represent vertical integration. So where you see M4 is connected to H2, H3 and H8, that means that these guys are all in one network. That's what-

PART 3 OF 5 ENDS [01:36:04]

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-sharing 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, F, Z, Z, Z, I don't know, that has two ways of looking at how preferential status and integration interact with where consumers eventually go. You can think about this as some evidence, some very, very preliminary raw evidence of steering through preferential tiering. So the left-hand table here shows you the percent. So on the X-axis, you have all the hospitals. Let's see if I can show you. Yeah, if we want to just keep some numbers in mind, let's take the first bucket, the first box. And remember, we have insurer MA that's vertically integrated with hospitals two, three, and eight.

So we can follow these a little bit in our table. So the left-hand table shows us the percent of plans from each insurer that give preferential treatment or have a preferential tiering for each hospital. So on the X-axis, we have the hospitals. On the Y-axis, we have the plans. If we look at the top row, that's insurer MA. We remember two, three, eight. We see that two, three, eight are shaded in darker blue because they have many more of their plans giving preferential tiering to the hospitals that are vertically integrated. It's worth noting that there are other hospitals that are not in their vertically integrated network that also have preferential tiering. It's just that the vertical integration corresponds strongly to having much more concentrated preferential treatment. We can see this even more so for insurer MB. I didn't emphasize it for you when we look back to the picture, but the first hospital associated with MB is hospital H4, which you can see is very, very dark blue.

Now, why does this matter? We can take another snapshot on the right-hand table, which looks at a similar kind of thing, but instead of looking at preferential tiering at the plan level, it looks at the percent of admissions. So what percent of people that are with insurer MA, let's say wind up being

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 can see this more directly. I had a hard time doing the comparison, so I just plotted them in the graph. So perhaps unsurprisingly, there seems to be a pretty strong positive relationship. Whenever hospitals are integrated, they're more integrated with an insurer, they're more likely to be preferred, and they're also more likely to get patients from that insurer.

Now, it's worth saying, going back here, that there aren't all that many zeros here. So all hospitals, I think, more or less admit patients sometimes from all plans. And so there is lots of shopping around. This is going to be some useful variation, but the main point here is that there seems to be pretty strong evidence of steering to specific hospitals based on preferential treatment at all and indirectly through the vertical integration. I'll show you a little bit more about that. Actually, I'll show you more about that right now. So let me show you a little bit more evidence that the steering is not just coincidental. It doesn't seem to be just that people choose the insurers that are integrated with the hospitals that they like, but also that there seems to be, as far as we can tell, some I think convincing evidence in the direction of the steering being somewhat causal.

One way to look at this, not perfect but goes in the right direction, is that if you look at people who move, and this is another figure from one of their appendices, people who move locations and therefore choose a new plan, and you look at the event study of how likely they are to wind up going to hospitals that are preferentially and that are vertically integrated with the insurer that they move to, and generally speaking, also preferentially tiered, you see that there is a sharp increase. The easiest way to see this is on the left-hand event study there. You see that up until times zero, times zero is when they move, there's basically no difference in how likely they are to use a preferentially tiered hospital. As soon as they move, they shoot up.

Okay. Now, why does this matter? So I mentioned that they use this notion of allocative moral hazard, and what they mean here is something somewhat subtle that Ben mentioned somewhere in his slides. So the general idea here is that because... Hospitals are specialized in different things, but many hospitals are good, even if they're not the best, and sometimes even the public hospital is quite good, maybe at simple procedures. The way that they look at this in their model or in their paper more generally is by looking at people's choices of hospitals for each diagnosis that they have.

So they use this diagnosis model, the demand for hospitals per diagnosis to look at the value which they define as a quality or hospital specialization for each hospital treatment pair. This is what this graph is showing. Now, it's important to note that the axis here, the Y-axis here goes below zero. And the reason is that the baseline value here is the public option. The public option, which is generally cheapest, I should mention. And what this shows you, if you stare at it really hard, there's lots of different colors, but what this should show you is that

many of the hospitals that are vertically integrated actually have blue spots, less red and even blue spots, which suggests that the public hospitals are better from the perspective of consumers at some of the treatment diagnoses.

The reason this is important is because what they mean by allocative inefficiency or allocative moral hazard here is that preferential steering through the vertical integration plan structure basically gets people to use the same hospital more or less all the time. And sometimes this is efficient because these hospitals are good at some things, but in many ways, many times it might be inefficient because the people are going to be using the same hospital for treatments that they could have gotten elsewhere, in ways that absent the specialized pricing structure might be inefficient from the consumer's perspective. And from the social perspective is inefficient in any case.

So the last thing, I'm out of time here, and so I just want to show one more thing. Basically, the model operates under the assumption that there's going to be lots of readjustment. Ben made a big point at the end of his talk that most of the action in terms of getting the welfare from removing vertical integration to be positive at the end is the fact that the formerly vertically integrated insurers are going to readjust their plan structure and remove the way that they do the preferential tiering of the hospitals that they're integrated with.

One reason to think that that's not crazy is that they can actually see an example of this in their data. Over the course of their panel, they have a vertically integrated pair that winds up disintegrating, and they can use that variation to look at the effect of VI on preferential tiering, and they see exactly what Ben's model is showing. They basically see that when a hospital insurer pair disintegrates, the probability that a formerly integrated hospital preferentially tiered goes down. Now, if you pile up all of the pieces of evidence that I've shown you before, you should expect that once that happens, you should see that prices change in the right direction, and that the consumers of the hospital insurer pair readjust as well, and this is exactly what their model shows. Basically what it shows is that without vertical integration, the formerly vertically integrated insurers will raise their premiums because they don't no longer internalize the value that the hospital is getting from the steering.

The formerly VI hospitals sometimes raise their prices and sometimes lower their hospitals, but most of what happens and most of where the action is is that basically the formerly vertically integrated insurers wind up changing their coverage and they increase coverage for other hospitals and decrease it for their own hospitals. And as a consequence, the consumers that were formerly being steered to using expensive hospitals for conditions that don't need expensive hospital care would go elsewhere, and in particular, would go to hospitals that are public or just more centrally located, which their demand model shows is very important. That's where most of the efficiency gains are coming from. Now I'm out of time and Ben already talked about the additional bells and whistles that they've done, and so let's leave it there. Thanks.

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-preferencing 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 self-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 8: Over here.

Benjamin Vatter: Behind the pillar. All right.

Speaker 8: I'm a big fan of the paper. It's incredible work. I'm excited to see how this goes going forward. My question is unpacking the consumer surplus gains from banning VI. There's this result in healthcare that consumers seem to value convenience way too much than seems reasonable compared to clinical quality. How much, if you unpack that, of the consumer welfare gains comes from convenience versus cost saving versus maybe something more like clinical quality? Do you have any sense from the estimates, what might be there and then what we might learn from that?

Benjamin Vatter: Yeah, I think that exactly what you're saying is pushing our surplus measures down. For a lot of people, we say they seem to be going to shopping to a hospital that has minimally different quality, but they've valued it a lot and they're really discounting the travel distance a ton, for example, for something that is actually meaningful and that is all exactly within this line is pushing our things down. If I put what I thought was the objective value of allocating the people, it would be much larger. So I think it's within this. In other parts of my work, I would've taken this fixed effect for the hospital diagnostic and tried to decompose it in the true value, the true quality, separated from beliefs. Here, we're just taking beliefs, we take it as given and those beliefs are going to be attenuated relative to the true gradient.

And so I think that it speaks to that. We don't have a direct decomposition that says, "This is the value of benefits." It's hard for us because it's hard for me to tell without putting my ideas on it. Say this is convenience versus this is truth, but it is definitely expressed there and my feeling is that it's depressing the consumer surplus gains.

Speaker 8: You're saying it would be even stronger without [inaudible 01:54:11]

Benjamin Vatter: Yeah, because-

Speaker 8: Oh, okay.

Benjamin Vatter: My feeling is that right now, we get a lot of misallocation. So this allocated moral hazard that [inaudible 01:54:17] discussed is a lot of people end up going to these hospitals that are really high quality for ingrown nail. And it's like why are you going to this extremely complex care for something that is so menial? But there's the convenience element to it, like they're putting a lot of value on the unobserved preference for that thing. And so just from terms of quality of care, they would be much better off going to say downtown hospital, simple, in and out, low complexity, low fixed cost of operation.

Speaker 5: Last question.

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-driven 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 in 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 under-enforcement. 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 downstream consequences of rising prices for mergers are very, very real. We're going to see job losses for workers in the lower and middle part of the income distribution outside the healthcare sector. And the third takeaway for us is that enforcement is actually cost-effective for the federal government. Because people are losing their jobs, tax revenue goes down, and the drop in tax revenue is actually much, much greater, we think, than the cost of taking an enforcement action.

So this number, \$25,572, this is really the number that motivates a lot of my work. It's premiums for health insurance for a family of four. Wild to think about. You can get a pretty nice car for \$26,000 because the modal individual in the US gets his or her health insurance from an employer. We're in a position where the average family in the US is buying a new car's worth, a new Toyota Corolla's worth of health insurance every year. So I study hospital pricing largely for the same reason Willie Sutton robbed banks. That's where the money is. It's the plurality of health spending in the US. It's a sector of the economy that accounts for 6% of GDP. A lot of my work has shown that the variation in health spending in the privately insured is being driven by variation in hospitals' prices. We know that price growth in the hospital sector is driving health spending. In fact, prices in the hospital sector have grown faster than prices in any other sector of the economy over the last two decades.

Now, this has happened concurrently with a significant amount of consolidation in the hospital industry, about 50 mergers per year. And if instead of graphing this in deals, I graphed it in deal size, we actually see it going up steadily over time because in a sense, a lot of the singleton mergers happened in the early 2000s, and now we're starting to see bigger systems come together.

What we've also seen is a modest number of enforcement actions, largely about, depending on how you count them, 13 over the last 20 or so years. And what we're interested in is, is that the right equilibrium? And we think a simple test to adjudicate this is whether or not we can see transactions that we can predict ex-ante, using standard screening tools that ex-post raise prices. And there are all sorts of reasons why we might think this is the efficient level of enforcement right now. Maybe the deals that are coming out and going through aren't problematic. Maybe they aren't lessening competition. Conversely, maybe the current enforcement rate is enough of a deterrence that the deals that are coming out aren't problematic.

So what we're going to do is take the two standard screening tools used by the agency. Right? The first outline in the horizontal merger guidelines, whether a deal raises the HHI of the merging parties by 200 points or more and leads to a post-merger HHI of 2,500 points or greater. We're going to use that as a flag. We're going to flag those transactions that cross that threshold. And second, we're going to follow some of the work that Devesh has done. We're going to build a willingness to pay measure for each transaction, which looks at where patients get care and then thinks about the demand in that sense that insurers have for including certain hospitals in their network and use that to predict the markups from a transaction. We're going to flag transactions that we think generate an increase in willingness to pay of 5% or more.

So I'm going to try to get through a lot. So I'm not going to do a deep dive into the methods, but what we do is we've built over the last couple of years, a database which we think is comprehensive of all mergers that have happened in the US hospital sector over the last 20 years. So we're going to rely on that. We're going to focus on mergers from 2010 to 2015. We're going to bring together a claims database from the Healthcare Cost Institute, which is going to have all claims from Aetna, Humana and United, but a third of the individuals in the US with private health insurance. And we're going to do a diff-and-diff. Diff-and-diff, we're going to take into account staggered timing, introduce hospital fixed effects, and then do matching. And we're going to see what happens to hospitals' prices post a merger.

Over time, we're seeing that about 20% to 25% of the transactions we see in the hospital sector could have been flagged as crossing the threshold that are in the 2010 guidelines. Right? Deals that really do substantially raise the HHI of the merging parties. And when we split these deals down into the price effects they have on mergers, what we see is that mergers that do raise concentration in a meaningful amount increased prices by 5% to 7%. Conversely, which I think it's also really important to point out, is that the 80% of transactions that don't lead to meaningful changes in HHI really don't have a major impact on prices. We also do this with the willingness to pay measure and we see very, very similar results. Mergers that we can predict as bad turn out to be bad. We take that as evidence that there is potentially under-enforcement.

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 pass-through into wages.

I think what's very, very different about the price increases generated by hospital mergers is it isn't creating a change in benefits that workers necessarily value. We don't see that these transactions raise quality. What they do is they make existing benefits more costly. And under those circumstances, the theory begins to shift. You might not think that workers are going to internalize that cost and be willing to give up wages. And as a result, we might see changes in the extensive margin, actually seeing job losses instead of wage cuts.

The second reason that I think it's so important is really the development of this recent literature that's posited that the way we finance healthcare in the US, employer-sponsored health insurance, is becoming and should be viewed as one of the leading drivers of income inequality...

PART 4 OF 5 ENDS [02:08:04]

Zach Cooper:

=should be viewed as one of the leading drivers of income inequality nationwide. And this is worked by Saez and Zucman, Amy Finkelstein, Owen Zidar, Anne Case, and Angus Deaton. And the basic idea is that wages vary a tremendous amount within firms, but insurance premiums do not. And so in practice, they look a whole lot like a head tax and we often think of head taxes as quite regressive. And so you can think of an insurance plan for ESI that costs on average \$6,000 per person. A firm can save \$12,000 by letting go of two

\$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 four-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-merger in our data was actually Yale New Haven's acquisition of St. Rayfield's Hospital. All politics in this one are local. That raises the cost for me of retaining RAs by raising health insurance premiums. I can't pay an RA less. All I can do is hire fewer Ras. In practice, I can't have half units of RAs. I'm letting potentially whole RAs go.

So what we want to do in this paper is establish this causal chain. Let me walk you through each set of our steps to make that happen. So we're going to build a panel of firms. We're going to focus on firms in the US inside and outside the healthcare sector that have at least 50 employees. We're going to do that using EINs in the tax data. And then we're going to link to all of the W-2s per EIN. So for each firm, we think in the US over 50 employees, we're going to measure the count of workers they have based on the number of W-2s we see, and then their total payroll, which is going to be the sum of the incomes on the W-2s.

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 spending onto labor market outcomes. So we're going to instrument that, and our instrument's going to hold health spending or healthcare quantities fixed. It's going to hold the providers we think an firm's employees fixed, and allowed the only thing that adjusts health spending over time to be that firm's exposure to the price increases from transactions.

So here's the intuition. Most of Yale's employees are located in New Haven County. We know from claims data that most of those folks get their hospital care at Yale-New Haven. We know that the Yale-New Haven merger raised prices, let's call it 25%. And so Yale's really exposed to the price increases from that transaction. The local power company in Connecticut has some of its workers in New Haven, but they also have workers across the state. So they're going to be in a sense, partially exposed to that transaction. The FTC or maybe Harvard is going to have some folks who actually live in New Haven. So some of their workers are going to be exposed to that merger because that's where they get care. But it's going to be a pretty small portion.

And we're going to do that estimation for every employer and every hospital merger across the US. And so in a sense have this measure of how exposed every employer was to the price increases of all the transactions that were happening nationwide based on where their patients get care or their employees get care. And the nature of the price increases across all transactions. And the key identifying assumption, what we're going to have to prove to you, is that those changes in hospital prices from these mergers aren't correlated with changes in firms' payroll or their employment. We're going to go to huge pains to show you that.

I think in some ways the most convincing is actually this panel on the right, and this is going to be with local labor market outcomes. We can do the same exercise with firms. We're going to show you the correlation between changes in county-level increases in prices from mergers and things like local spending per beneficiary, percent changes in income per capita. This is in the pre-period. And the percent change in the share unemployed. And what you see is the scale of the price increase induced by transactions is uncorrelated with our pre-treatment trends in local economic conditions. We see the same thing at firms.

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-wage 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 non-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 in prices lowers local labor income by a \$1.33. So it turns out there is a deadweight loss from these hospital mergers. It's just happening in the adjacent labor market. Because demand elasticity for hospitals is pretty low and quantities don't go down. The reduction in quantity is actually happening among workers. And I think that's really, really important to point out. Now there's a literature which in some ways I really wasn't familiar with, but I should have been. And it's that the downstream health consequences of losing your job are

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-harm 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-diff point estimates. Now, if you thought that the diff-diff point estimates are way too low, it's just going to change that coefficient, it's going to mute down the

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: We're not being conservative.

Zach Cooper: I don't think we are.

Speaker 10: Hi Zach. Really provocative arguments here. If I remember correctly when I was working at FTC, it's not just economists or lawyers in FTC don't want to crack down on hospital mergers. It's also not even resource problem that we don't have people working on those. But in some cases, state or some local laws, they can give exemption to the enforcement's action from federal agencies like FTC. So I wonder whether you thought about that as kind of another way to change the enforcement action?

Zach Cooper: Yeah, it's a great question. So I just want to be really clear. I don't think this is a question of lack of will from the agency or visibility into these transactions. I think what you've raised is a really good point about states' willingness to go along. I think there's a really good example of it right now. There's a straight two-to-one merger right now in Terre Haute Indiana. The state has come out for what's called a certificate of public advantage, which in a sense shields the transaction itself from federal investigation. The FTC issued a public comment, which I thought was wonderful, that said, depending on how you measure it's going to raise the HHI locally from 5,000 to 10,000. That's not good. That's likely a problematic merger. Our estimates are this transaction's probably going to lead to price increases of 20%. We can use our simulations to say our guess is about 800 job losses.

I think it's two things. So one is an implication of the work and one is talking about the work with local communities. So I think the first is this idea that when these mergers happen, first we don't see clear evidence that they increase employment at the merging parties, but more generally they lead to widespread decreases employment overall. So it's leading to net reductions in employment. I think a lot of what folks think about at the state or local level is, well, these are big pillars of the economy. When they come together, it's going to be good. And I think our evidence just doesn't suggest that's the case. And I think part of, and this is the Tobin mission, but I think it's true for all of us, is how do we have those conversations locally? It's getting the subtlety out there. It's talking to state AGs and local policymakers and saying, "Look, we know that this is hard and we know you're under enormous pressure, but here's what our best evidence about these transactions can suggest."

Speaker 11: Do you think there's anything from your results that we can learn about the effect of quality from mergers? So for example, you show the deaths from despair from the unemployment, but you are not finding results for other types of diseases such as cancer, things like that. Should we learn anything from that about the quality effects of mergers?

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]