[SQUEAKING]

[RUSTLING]

[CLICKING]

TOBIAS SALZ: OK. Thanks so much for being here. I know it's a very busy time of the semester. So I'm going to talk a little bit more about search models. And I know that Glenn has already introduced you both to the theory side and some empirics. So I'm going to be brief and motivating the topic.

So there was a famous article published by George Stigler in 1961 where he made the observation that in many markets, it's not the case that there's a central marketplace where everybody observes the relevant prices. But oftentimes you actually have to move around, go to different places, incur physical search costs to figure out what prices are. Or maybe it's mentally costly to figure out what the right price is.

And so that introduces frictions, which he termed search. And so in this quote here, he says that we can look at the dispersion in prices. And in some sense, they're a measure of the amount of ignorance in the market.

So why do we want to look at economic settings through the lens of a search model? Well, so first of all, there's several positive arguments why we may want to do that. So first of all, in many settings, we do see that products that are very similar or essentially identical are sold for very different prices. And so as researchers, we would like to make sense of that observation how can this be.

And again, so one potential explanation which we will explore today empirically is that with search frictions, it could be that you get a mixed strategy pricing equilibrium, which leads to equilibrium price dispersion. But then within the theme of empirical IO, it's also normatively important. So when we quantify, let's say, welfare in markets, we might want to take search frictions into account foremost because they're actually a source of market power. So in IO, we are predominantly concerned with various ways that firms gain market power. Such frictions are one potential source of market power.

There are also, depending on how you think about them, a social cost. So if people have to engage in wasteful effort to drive around to various car dealerships to figure out what the right car prices are, that's socially costly. And if we do welfare calculations, we may perhaps want to account for those costs.

Such costs can also lead to misallocation. So if you think about you're looking for, let's say, a hotel for your vacation or, let's say, a car or something else, you might not, in fact, find the product that you would like to purchase the most at a given price point. And similarly, if you think about a firm, let's say, purchasing inputs from a supplier, the firm might not find a supplier that can offer the specific product at the lowest cost. And so that leads to welfare losses that, again, we might want to quantify.

So this is general. For the second paper that I'm going to discuss today, I will look at markets in which prices are not posted but negotiated. And these are very interesting markets. And we have recently much more data which allow us to make some empirical progress on these kinds of markets.

So if you think in many settings, people negotiate prices with the seller. So if you buy a car, a house, a mortgage, so many kinds of consumer financial markets have this feature that prices are somehow negotiated and various kinds of services. So if you get a contractor for your home or, let's say, a plumber or something like that.

So in these markets, it's oftentimes particularly hard for consumers to know what is the fair or the right price, because searching is particularly costly and you have to go through some process. And there is oftentimes this element of idiosyncratic pricing that makes that even more difficult. What makes this also very interesting, and this is something that we will explore today, is that there are potential distributional implications of pricing And policies such as mergers that we think about as economists.

An empirical challenge that arises when looking at those markets is that if sellers are quoting different prices to buyers, but we only observe the transaction prices, this is a selected set of prices. And so we need to know how we can map this back to the prices that are actually offered.

So here's the roadmap for today. I'm going to talk first about a paper that's kind of conceptually very interesting and ties back to the theory lecture that Glenn held on search. And then a more empirically oriented applications on mergers and search markets. And if I have time, I'll probably not have too much time, I'm also going to talk about my job market paper for a little bit, which is in this space.

So let me talk about this paper by Hong and Shum. So it's really an interesting exercise. And one thing that makes it interesting is it takes really this idea that was in Stigler's quote to the extreme. So it asks what can we learn from a price distribution? And it imposes a specific equilibrium, namely that firms are playing mixed strategy, mixed strategy pricing. The model that it builds on is Burdett and Judd. This is conceptually quite close to the style model that you had seen in Glenn's lecture.

This data only-- this paper only uses data on prices. No quantity data. And so, again, the idea here is we take these mixed strategy pricing equilibria very, very seriously and see how far this gets us empirically. And at the end, we can discuss a little bit whether this is a sensible way to go about this.

The other interesting thing that this paper does, it looks at basically the same problem with two different models. A sequential and a non-sequential search model. And it turns out that empirically, it matters which one you use. And I think that's a good exercise to go through, to think a little bit about you're conceptually interested in something. There are different kinds of models, slight variations that might allow you to in some cases recover primitives better than in other cases. And so I think in that, the paper is quite nice that it helps you sort through these things.

So in this setting, consumers are taking non-sequent-- so in the first version of the paper, non-sequential draws from the price distribution. So they're basically committing ex ante to l number of draws. And the way they do this is by minimizing their total cost, which is their search cost. Each consumer has an idiosyncratic cost for searching that's constant across searches. And then there's an expected price that results from taking l draws from the price distribution. And one of the draws is for free.

So one thing I should have mentioned is that consumers in this market, they have a unit inelastic demand. That means they purchase one product for sure, but not more than that.

And so this here is the price distribution that comes out of the mixed strategy pricing equilibrium. So again, this expression here is just the expected lowest price out of taking l draws. Any questions about what consumers are doing here?

OK, so how can we think about this? So consumers are comparing the marginal cost of searching, which is ci, to the marginal benefit of searching. What is the marginal benefit of searching? Well, if I think about going from i to i plus 1 draws, it's how much lower the price is if I take an extra draw. So the marginal benefit of searching is the difference in expected prices from taking another draw from the price distribution. And so they're comparing marginal cost and marginal benefit. As long as the marginal benefit is higher than the marginal cost, the price reduction, they're going to increase the number of draws by 1.

It's going to be helpful to keep track of these types here, the delta i's, because these are consumers. Consumers, again, have heterogeneous search costs who are just indifferent between drawing i and i plus 1 times from the price distribution. So we get this kind of sequence of marginal benefits.

Now, with this in mind, there's actually a really nice way of displaying what consumers are doing in this market. So this here shows the CDF of-- sorry, the density of the search cost distribution. And consumers are ordered in terms of how many draws they're taking from the price distribution. Consumers with a very high search cost, with a very high search cost, they only take their free draw. And as you lower search costs that take more and more draws from the price distribution. And again, these are the difference types that segment these different bins of the distribution of search costs.

So we get these kind of mass points of consumers that take i draws from the distribution. And we call those q tilde. Any questions on this? In this direction, search costs are increasing. With lower search costs, you're taking more draws from the price distribution.

OK, so let's now think about what firms are doing. So this is, again, based on a model by Burdett and Judd from 1983. Firms are posting prices here. And they're doing so maximizing the margin times the probability of selling to a consumer. These q tildes, remember, were how many consumers are taking k draws from the price distribution.

So the q tilde k of them, they take k draws. And I'm selling to them if I'm offering them the lowest price. And so, again, you have to think of this here as the price distribution that comes out of the mixed strategy pricing equilibrium. This is a smooth, atomless distribution has some support points that will be important in a second.

So how can we get some traction on this empirically? The restriction that we're going to use here is the fact that firms have to be indifferent between offering all the prices that are in the support of the price distribution. As is the standard condition in a mixed strategy pricing equilibrium. Everything that's in the support, you have to be indifferent between those strategies. And so in particular, firms have to be indifferent between offering any price p and offering the highest price and only sell to consumers that are not searching.

So now let's stare at this for a second. So what are the things that we do know and what are the things that we do not know in this equation? We do know the prices. Again, as I told you earlier, they're starting from the assumption that the researcher observes prices in the market. But we do not know how many consumers are searching how many times, as well as what the marginal cost is of firms.

So this gives us conditions. Linear equations, in fact, that allow us to figure out what q and r are. Again, because we observe prices, we can construct this CDF of prices here.

So as long as we observe more prices than the maximum number of draws that consumers can take in this market, or we have to observe more prices than k plus 1 of that, we can build moment conditions or equations from which we simply back out these unknown q's and r. So this just amounts to basically counting the number of equations relative to the number of unknowns. So these are sets of linear equations.

The way this actually works in the paper is that you typically are overidentified because n minus 1 is much larger than k, and so you use, let's say, a moment based method or, in fact, in the paper, they use empirical likelihood to back out these unknowns. Again, so what are the unknowns? I have k minus 1 of these unknown q tildes. Remember, this was this partition of the search cost distribution. And you have the unobserved marginal cost of firms, which is r, and assumed to be constant across firms.

STUDENT: So just real quick, a place where Hong and Shum are cheating a little bit is they've got that f hat of pi, the CDF evaluated at pi, and they're setting that CDF evaluated pi in the empirical CDF evaluated at pi. So it's as if-- it's not the prices are drawn from a distribution. They're assuming that is the true CDF at every observed price. The number of prices is less than the pi. And that's where their equations come from. Because at every observed price, they assume they know f [INAUDIBLE].

TOBIAS SALZ: Yeah, I'll come back to this because if you, if you actually look at the price distributions, they look a little bit funny. They have mass points. And so it doesn't look super consistent with an atomless kind of smooth price distribution.

So the bottom line here is that as long as we have more observed prices than k, and they haven't really told us where k is coming from. I'm going to talk about this later on. But as long as n minus 1 is larger than the maximum number of draws that consumers can take, you can use these equations here and simply solve a system of equations for these unknowns. Then again, practically, this is going to boil down to a moment condition.

OK, so let's now move to the sequential search model. The sequential search model is using a solution that you might have seen from, let's say, the McCall model. I'm assuming you have seen the labor search McCall model in macro or somewhere. But let me remind you how this works.

So instead of committing ex ante to a certain number of draws from the price distribution, you're going to draw a price and decide whether or not to draw again. And how do you make this decision? Well, the typical solution to this kind of problem is a cut off rule. So there's a reservation price. As long as the prices that you're drawing from the price distribution are higher than the reservation price, you keep on searching. But once it falls below that, you stop.

And so now the question is, how do you actually set this reservation value? And so here, again, the same idea as before. You're thinking about what is the marginal cost of searching. Again, this is constant and equal to ci. And what is the marginal benefit of searching? Which is basically the weighted average of price improvements that I could get where z is my reservation value. Again, this is implicitly a function of my search cost.

And so you get this weighted average of price improvements. This here is the marginal benefit of searching and ci is the marginal cost. So z is set such that these two are equal to each other. Does everybody-- may I ask, have you seen the McCall model? So some of you.

So with this solution, it turns out that empirically, you can make much less progress, and in fact, not rely on the same argument as we did before. So why is that the case? So first, let's think about the reservation prices.

So one thing that's different here is that beforehand, we had this sort of discrete number of things that consumers could do. So they were taking a continuous draw from the search cost distribution. But still, in the end, there are only finitely many things that they could do. So you get these kind of bins of the search cost distribution. And we were identifying these bins.

Here what is different is that the optimal strategy is a function of ci and every unique search cost type has a unique reservation value strategy. Except for those that potentially bunch at the top. So if you have a higher reservation value, then the highest possible price in the support of the price distribution, then of course, you stop searching when you reach-- when you drop p bar.

So now let's think about the demand of firms with this kind of strategy. So as before, we can construct indifference conditions. So what is the demand of a firm? I'm going to sell only to consumers that have a reservation value higher than the price p that I'm setting. So the demand is proportional to the mass of consumers above the price p that I include, that I charge.

So now again, we can try and construct these indifference conditions for any given price pi. Firms need to be indifferent between charging that price, reaching consumers only above that have a reservation value above that price and the maximum possible price in the support of the price distribution. Can you see what the issue here is?

So the problem is we have a incidental parameter problem. For any nu of those indifference conditions that are introduced, I'm also introducing a nu unknown, which is the mass of consumers above that point. So in other words, this simple moment counting versus unknowns no longer works because I have n plus 1 unknowns. The mass of consumers below pi and also the unknown cost of the firm. But I have n minus 1 indifference conditions.

So basically we can no longer use this argument. So in the paper, they say to non-parametrically identify the search cost distribution, in the previous case, we also didn't quite non-parametrically identify the search cost distribution because we only got this partition into these bins. But here we basically with this strategy can learn anything about the search cost distribution unless we're willing to impose some parametric form. And so that's what they're going to do in the end. They basically assume that the search cost distribution takes on a specific parametric distribution and then simply estimate it based off of that.

So this comes back to Glenn's point. What is the application here? The application is perhaps not the most inspiring, but maybe in this crowd, which is the different sets of econ graduate textbooks. So Stokey-Lucas, Lazear, Billingsley-- *Probability and Measure.* So if you stare at these price distributions, you can ask yourself whether you think it's plausible that these are coming out of this mixed strategy pricing equilibrium.

So one thing that I, for example, find somewhat curious is that you get these very large mass points here. And remember from Stahl but also from Burdett and Judd, we would expect an atomless smooth distribution of prices.

Let me show you some of the results. So first, they have to figure out what k is. Remember, k was the number of draws from the maximum number of draws that consumers take from the price distribution. The way they do this is somewhat hacky. They basically just look at what is the highest k for which my model converges. So estimate the model with that.

And so one striking thing that they find is that for most of these different markets, the average consumer only searches once. So about 50% of consumers only take one draw from the price distribution. And in fact, only for one of the textbooks, Lazear, you actually have consumers that take three draws from the price distribution.

So if you think about this here, one potential issue with having 50% of consumers only searching once, which remember you get these cut off points for consumers. Let me just go back to this briefly to illustrate this point. So what I'm identifying here are these partitions. So I'm getting out of the model a partition of the search cost distribution.

Now, if let's say delta 1 lies very, very low and most consumers only search once, then I don't actually learn all that much about the price distribution above that point-- about the search cost distribution above that point. So the problem is that what we're really recovering here is not a primitive but an endogenous object, which is the search strategy of consumers.

And so we have to, in some sense, cross our fingers and hope that how consumers are searching in the market tells us enough about the search cost distribution. But it could just be that most consumers don't search at all. And then I don't really learn anything. So that's kind of one potential shortcoming of this approach is that this identification argument doesn't directly map to the distribution, but it maps to what consumers are doing in the market.

If you think about this paper and let me comment a little bit on this, if you think about the exercise, just using price data is a somewhat heroic exercise. And you have to think about this as being this was one of the very first empirical papers on search or papers that use the equilibrium conditions that you get out of a search model to say something empirically.

So one thing you may want to ask yourself is how plausible you find it that firms are really setting prices by mixing. And whether there are alternative explanations why we see different prices. This wouldn't be on the top of my list of complaints, but the marginal cost of firms are assumed to be constant here. So you could take some issue with that.

And then I want to point out something that is more of a common critique of many papers in the search literature. So first, most papers find very large search costs. And one issue is that you look at the dispersion of prices and you say that all of this dispersion, all those differences across prices is really kind of structural variation in a sense that all of that variation has an economic interpretation.

As opposed to things that happen with data, it could be that there's, let's say, measurement error in prices or other reasons why the prices that you observe are not exactly equal to each other. And so in these models, you load typically almost the entire variation of prices onto this search cost distribution.

And so that's part of the reason why in many of these papers, you find somewhat excessively large search costs. And so I think that's something that this literature has to grapple with a little bit. But it's actually common to other literatures that interpret all price variation as structural, which is, for example, also true in auctions.

So the other thing that you could potentially complain about here is that these sellers of textbooks, they might be differentiated in other ways. Some of them might have funny names or located in places that people, I don't know, don't like that much. And so for that reason, they have to charge different prices for the same kind of product. So although these are in principle very homogeneous products, you could think there are other unobserved differentiating factors that lead to differences in prices.

So again, this is the way you have to think about this paper. This is, I would say, mostly kind of a conceptually interesting exercise. It teaches us something about using these equilibrium restrictions. It's perhaps not a paper that you would use nowadays to really think about such frictions, in part because of this very strong equilibrium restriction that firms are playing mixed strategies, which in some settings we know is sensible. Let's think about, for instance, sports and other places where we know that mixing makes sense. But here, you can do some introspection. I find it somewhat less plausible.

So what I'm going to talk about next is a paper that is much more empirically motivated by Jason Allen, Rob Clarke, and Jean-Francois Houde. It came out in the AR in 2013. And it basically asks, what are the implications of a merger in a market where prices are negotiated? So just to remind you a little bit, so in IO, we have a very developed machinery for thinking about mergers in posted price markets.

So we estimate demand. We have some ownership matrix with demand. And the ownership matrix we can compute first order conditions and can recover costs. And once we have costs, we can, let's say, change the ownership matrix and recompute prices. Up until very recently, there was not a lot of work that thinks about what are some of the implications of mergers in markets where different consumers pay different prices.

So they're going to look at this in the Canadian mortgage market. They have some great data on insured mortgage contracts, which are collected directly from the insurer. And what is nice about the fact that these mortgages are insured is that they can claim that these contracts are very, very homogeneous, that consumer risk is not playing a large role in explaining price differentiation. So they claim that this is a very, very homogeneous product because idiosyncratic risk here is insured away.

So part of why I like this paper is it provides a very nice combination of treatment effect analysis of an actual merger that happened in the data and interpreting this variation through a structural model. And so I'll try to make this clear today, but there's basically a very nice connection between the treatment effect analysis and the structural analysis.

So there's a merger between two large lenders that happens in the data, which has implications for different markets. In some markets, you have branches from both lenders. In some other markets, none of the lenders have branches. And then that changes the choice set of consumers differentially across markets.

So here's a picture of what happened in this market over time. So one thing that's pretty clear and visible is that over time, this market has become much, much more consolidated. So initially, there were banks and credit unions, trust shares. So these are basically lenders that are owned by their members. And then there is what they refer to as the big eight. These are the largest banks such as TD Bank, Nova Scotia, The Royal Bank and others. And these have been subsequently overtaking trusts. There were also some regulatory changes, and so the market has consolidated over time.

So some institutional details that are going to be relevant. So consumers here have to go to a lender to figure out what the price of their mortgage is. So they have to talk to the local branch managers. And the local branches, they're posting prices, but the consumers can negotiate discounts from those prices. And they have explicit incentives to not give too many discounts, because their compensation is affected by the level of discounts. And as they claim in the paper, they also are explicitly told not to compete with other branches that are owned by the same lender.

So because the data here has essentially administrative data that comes from these insurers, they see very, very rich information about both the consumers and the type of contract that they get. So they see, for instance, the consumer's home location, the loan size, house price, things like loan to value, and demographics such as income and even whether or not they had a prior relationship with this lender.

Another nice thing that they can do is because they see the consumer's home location, they can actually define markets around consumer's home address. And that's in terms of market definition usually pretty good, because then you don't get these weird cases where somebody lives at the border of a market and would potentially shop from a market that's defined as a separate one in your data. And so because the Canadians do this in kilometers, so of course, most sensible.

So what are the treatment and control groups here? So they're going to only focus on the short run effect of the merger. So that means that they're not going to look at entry and exit, but only at the immediate change in prices that the year after the merger happened. And so I'm only talking about the merger here because I cannot actually disclose who the lenders were that merged. So these were two larger lenders.

And so what are what are the control markets? There are some markets in which neither of the lenders or only one of them were present. These are the control markets. And then there are markets in which both of the merging lenders had a branch. And in those markets, competition is now reduced. Because, again, according to this story that lenders are internalizing and not competing against branches in their own network. The choice set has now been reduced. So consumers cannot compare as many prices as before. Any questions about the institutional details here?

So here's what I do. I first have a descriptive exercise where they look at the price variation that they observe in the data. And again, what's important here is that they can credibly claim that most of that is not because of risk variation across consumers.

So they first run a regression of the margin of a consumer, that's the rate that they observe in the data minus the bond rate, on a set of control variables. So how long the mortgage was, how large the size of the mortgage was. So the idea here of this regression is that you really want to get at the unobserved component of the price.

So there could still be differences across consumers, because again, one consumer gets a larger mortgage than another. Or there could be differences in FICO scores. So there could be some sort of perceived residual risk of a consumer. But they want to look at once we control for all these things, how much price variation is left after that? And so the key object of interest here is this what they call negotiated margin. It's the residual from this regression. So this is what they mostly analyze. So again, control variables here also account for time variation and so on.

So here are some summary tables. There's a lot going on, but I want to point out a few things. So up here we have control markets. And down here we have treatment markets. And so both before and after. So one thing that's true is that the control markets have somewhat different consumers than treatment markets.

So for instance, in the control markets, consumers' house value is somewhat lower. Consumers also have somewhat lower income than in the treatment markets. But on other variables, they're fairly balanced. So for instance, loan to value is almost identical. FICO scores are quite similar. The number of renters is pretty similar and so on. So they appear to be quite balanced on observables.

Again, the key object of interest is this negotiated margin here. And the initial negotiated margin is very similar across both markets. So this is here. Again, control and treatment in the before period. And it's 106 and 107 basis points respectively.

So now we can look at this negotiated margin in the after period. And in both markets, it goes up to 167 in the control markets and a bit more than that to 172 in the treatment markets, where both lenders, both merging lenders had a branch, and therefore the choice set was reduced.

So one thing that is interesting, you see it directly in the summary statistics here is the following, which is that if you look at this was the average effect that I just pointed out in the negotiated margin. But if you look at across the distribution of prices, you see that the 75th percentile of the negotiated margin is again very, very close to each other. They're very close in the after period. But there's larger difference, about eight basis points, in the 25th percentile.

So it appears that most of this effect of higher average negotiated margin is coming from the fact that consumers at the low end of the price distribution are now paying higher prices. And so this is kind of a theme that will run through this. I told you at the beginning that what makes this paper interesting is that they want to look at the distributional effects of mergers.

And it's tying this back to the theme of this class, which are search costs. The idea is that different consumers search differently as a function of their search costs. And what this implies for their search strategy, they get different kinds of prices. And if I'm a consumer who basically never searches, then the merger will probably not have as big of an effect on me. If I'm a consumer that searches a lot, then having an extra lender that I can go to is going to potentially affect my price. And so what is nice that is this idea is already visible here directly in the summary statistics.

So this would normally be now a great table to ask my 14.03 undergrads to compute the difference in difference treatment effect from this. But I'm going to just show it to you. So it's not very large. Depending on how you look at it, whether you account for trends or not, you can do this either with a matching estimator or not, but the results are fairly consistent. It's about seven basis points.

And so I think in the paper, they say for the average mortgage in this data, this implies about an additional $5 of extra payments per month. So the average effect is not large. And it suggests that this is a fairly, to begin with, a fairly competitive market.

But what it tells us is that before IO people started looking at these markets, one view of the world could be that most of the variation that we see in these kinds of markets is due to risk or cost differences. And what this shows is that potentially there's a market power explanation here. I think I said this. And again, for people who search a lot, you would expect this to be a fairly competitive market.

So this was the average treatment effect. So next we're going to look at the treatment effects at different points of the price distribution. So for that I have to introduce some additional notation here. So again, we're interested in the distribution of negotiated margins. And they're going to index those here with g and t, where g stands for whether or not you're part of the treatment or the control market. And t stands for whether or not it was before or after the merger happened.

What we want to make this comparison is a counterfactual distribution of negotiated margins for the treated markets. So we would like to know what, just as you would in your typical difference. In difference analysis, you make some assumptions to know what would the counterfactual market have looked like. So similarly here, we want to construct a full counterfactual distribution of prices for markets that were treated for the counterfactual where the merger didn't happen.

So what they're going to use is something that Athey and Imbens proposed in 2006. It's called a changes-in-changes estimator. It's basically a generalization of difference in differences where you basically look at treatment effects. We can look at treatment effects at different points of the distribution.

So the way this works, always have to write it like this to remind myself. So the way this works is we're asking take some margin and look at basically remember here t stands for before and after. g stands for whether or not this is a treated market. So we want to look for those markets that were not treated. How large was the change in the margin at a specific point of the distribution?

And we basically want that before for margin that gives rise to this across both markets is the same, because that defines the counterfactual distribution. So across all points of the distribution, we want to map from prices after the merger has happened to a set of prices before the merger. And under the counterfactual distribution, those have to be the same.

STUDENT: Why is the first subscript a 0, 1 and not a 1, 0?

TOBIAS SALZ: Let me see. So this here is-- yeah, this has to be 1, 0. Good catch. I apologize for that. OK. Sorry about that.

So again, we're looking on both sides here at the change. And again, the counterfactual distribution is defined such that from any point that I look in the after distribution, the margin that maps to this point has to be the same.

So now I can use this condition here. This is the object that I do not know, which is the counterfactual distribution of prices. And I can apply, basically, and yeah, here it's correct. So I can apply f 1, 0 to both sides. And I back out the counterfactual distribution of prices.

So what is nice about this is once I have this object here, which is, again, constructed from three observed distributions, I can look at different points of the distribution how are consumers treated. And in particular, I can come back to this idea that consumers that I expect to search a lot should be hurt more by the merger than consumers who do not. Any questions on this idea here?

So again, conceptually, it's like difference in differences, but you look at different points of the distribution. So here are the results that we get out of this. Again, you can look at those either what they call baseline or adjusting for trends. And what you find here is that if you look at the top end of the distribution, the confidence bands, they cross 0. So these are all not significant. But at the lower end of the price distribution, you get significant effects at the order of 7 to 10 basis points.

So again, what is the interpretation here? There's some consumers in the market that are searching a lot. And once their choice set is reduced, their prices go up. But consumers who were not shopping around to begin with, they're unaffected by this price change.

So this leads to dispersion, to reduction and dispersion and prices. So different measures, different dispersion measures such as standard deviation, coefficient of variation, and interquartile range all tell a consistent story that prices are getting compressed once the merger happens.

What they also do is they are looking at whether or not this counterfactual distribution of prices that they back out is equal or not to the actual price distribution that they observe. And they can reject this with a Kolmogorov-Smirnov test. So basically, looking over the full distribution of prices, reject the null hypothesis that the price distribution didn't change relative to the counterfactual price distribution.

So I think this is nice because, again, it tells a story that's quite consistent with the idea of different consumers faring very different in these markets. And you can observe it in those heterogeneous treatment effects.

So they now construct a search model to make sense of those results. And the model is constructed so that in the end, it maps back to these negotiated margins in the after period. So we have n plus 1 lenders in this market. And there are three stages to the model. So initially, as in Hong and Shum, consumers get a free price quote.

So you can think of this maybe from their home lender. I get a free quote. I can either accept this quote or I can reject the quote and engage in costly effort to gain additional quotes. And so they're using a search technology that's somewhat different from what you have seen before. So basically, there's a continuous effort function.

And as I increase my continuous effort, I get a higher probability that I'm getting more quotes from the price distribution. And the way that they construct this is that with probability s, I either get quotes from all the firms in the market and with probability 1 minus s, I only get two quotes.

So this turns out to be quite convenient, because they can basically map this just in continuous effort problem instead of dealing with something like a reservation price strategy in the sequential search model or having to compare different bins and where consumers sort themselves into. So this is somewhat similar to what people do in labor search. So I'll come back to that comparison later on. And it just makes the model slightly easier to handle.

So these are the three stages. Consumers get an initial quote. They either search or not. And when they search, consumer, the lenders compete for their business.

So let's work backwards. Let's first go to stage three. So this is where consumers went out and get additional quotes from the market. They modeled this as an English auction. Again, this is a fairly strong assumption, but it makes dealing with the model somewhat easier. So lenders here have a common cost component c and an idiosyncratic component epsilon j. And as you know, there is a weakly dominant strategy in English auctions, which is to bid your cost. So the lowest cost lender wins and is paid the price of the second lowest cost lender.

So with this, they can then define an expected price, or an expected negotiated margin, which is, again, a function of this common part here, as well as the expected second order statistic out of taking n draws from the distribution.

And since they make this assumption that consumers either take two draws or n draws, depending on whether or not their search is successful, you get this difference in expected prices or delta n as a function of how many lenders there are in the market.

So thinking about this, you can already see how this maps back to the empirics. So in some markets, the total number of firms that you can search from is reduced by 1. And of course, under this specific search technology, some consumers who would have taken usually n plus 1 draws are now only taking n draws when their search is successful.

So this was the last stage where the lenders are competing against each other. We now move to stage two where consumers are deciding how much to search. And so their total search cost here, which they call kappa i, is a function of how much effort they exert. And in particular, it depends on an idiosyncratic cost for each consumer UI.

So this is the consumer's marginal cost of searching, exerting more effort, plus a fixed component eta i. So if you reject your initial quote and you go out to search, you pay this fixed cost. It will turn out, I'll come back to this, that they basically dropped this in the empirical application, but I'll leave it in here for now.

So an important assumption and restriction in this model is that they assume that the marginal cost of searching of the consumer is publicly observed. So in other words, the lender, once you talk to the lender, you walk into the branch, they can actually somehow learn enough about you to know exactly what your cost of searching is. And this has quite strong implications. I'll come back to this. But this is something important to keep in mind here. So the marginal cost of searching is publicly observed.

They assume that there's matching probability. So the chance that the search is successful and I get n draws from the price distribution has this Pareto distribution here. And so with that, consumers compute their reservation value, which is the minimizer of the search cost plus the expected price that I get.

So again, this here is a common cost component. If I'm not successful, I'm getting the second order statistics out of taking two draws. And when I'm successful, I'm replacing this with the second-order statistic out of n draws. This is this delta n that we defined beforehand. The difference in expectations.

So it will be useful to keep track of this reservation value as a part that is common to all consumers and this idiosyncratic component here, which depends on the consumer's idiosyncratic search cost, as well as the number of lenders in the market. So in other words, consumers are different in terms of how costly it is to search for them. They exert in this continuous effort. This gives them some reservation value comes out of this optimization problem. And there are some consumers who will not search at all if their search cost is too high.

So we are now moving to stage one, which is the lender who gives the initial quote decides what that quote should be. So for that, we define the distribution of eta, which was this fixed cost of searching. So the consumer who decides whether or not to take more draws privately observes this fixed cost.

And the probability that the consumer goes out to take more draws is the probability that eta i plus the reservation value is lower than the initial quote that I get. The lender observes ui, so the lender knows what your reservation value is. But the lender does not know eta i. So this is privately observed.

So the lender who sets the initial quote maximizes, again, the margin times the probability that the consumer does not walk away. And there's an implicit assumption here that when the consumer walks away, then I'm losing the business of the consumer for sure. So the consumer then takes one of the offers from the other lenders and doesn't compare this to my initial m0 quote.

So this is basically a monopoly pricing problem, where I'm pricing against this residual demand curve here, which is whether or not the consumer walks away or not. And you get this solution here where the margin is equal to the inverse hazard rate. So a pretty standard solution for these problems where consumers have private information. So the consumer earns some information rent, and that leads in principle to somewhat lower prices than the reservation value would imply.

In the empirical application, however, they're going to just set eta i to 0. So the empirical application, they assume that consumers do not earn any information rent and lenders are therefore just paying them their reservation value, which is because they observe ui publicly observed.

OK, so with this solution, where there is basically now, according to this very stark restriction, no more any private information of consumers, you can directly map what comes out of this search problem of the consumer to the observed distribution of negotiated margins in the after period. So remember, this here is just for the treated markets in the after period. The fraction of consumers who get a negotiated margin less than m.

And we can write this as the probability that this reservation value, which had the common component and the consumer idiosyncratic component, is less than m. And if we define hu to be the distribution of these idiosyncratic search costs, we can instead rewrite this as follows. So we can write the quantiles of the consumer types, which is how costly it is to search for. You have to correspond to the quantiles of the negotiated margin.

So this is a strategy that you will also see in auctions where there's something kind of idiosyncratic that's unobserved. And consumers are ordered in some way. And that's going to map to something observed that's kind of ordered. In this case, the price that consumers get. And again, the way this works here is I have my search costs. The lender observes that and just pays me off my reservation value. And so you get this kind of perfect ordering of the negotiated margin that corresponds to the ordering of these unobserved types.

Of course, this here only works for a given set of parameters of the remaining problem here, which still includes the firm's cost variation as well as lambda i, which was determining how costly it is for consumers to search. So we can't quite directly back out this distribution hu, because we don't know yet the parameters, the remaining parameters of this problem, which are the search effort parameters as well as the cost of lenders.

So one thing that you get out of this is that consumers that have higher search costs benefit less from the merger than consumers with lower search costs. So for u prime larger u, you get a difference in reservation values for less sophisticated consumers that's smaller than for consumers that have a lower search cost. So in other words, what this model here generates is this difference that we already saw in the data that we see larger treatment effects at the lower end of the price distribution, where we would expect to be the consumers that have lower search costs relative to at the high of the price distribution.

So in other words, this pricing function under symmetric information basically is monotonically increasing in ui. It's weakly decreasing in n. So competition helps. And there are these increasing marginal effects of the merger when you get to consumers of lower search costs.

So with all of this put together, what they're doing is they now want to explain the actual observed differences in the treatment effects. So we had basically large treatment effects at the bottom of the price distribution, small treatment effects at the top. And they're going to find parameters, which is the common cost of lenders, the consumer search strategy, as well as the variation of the idiosyncratic cost of lenders, that best explain the differences in negotiated margins going from before to after the merger.

So this is, I think, quite nice because they basically let the model be informed by this treatment effect that they observe in the data. And so because of admittedly a fairly stark assumption, they get this direct mapping between the observed negotiated margin and what is unobserved to the researcher, which is the type distribution of consumers. They can impose basically that quantiles of consumer types have to correspond to this object that they get out of the model. So they're basically directly imposing the fact that these observed margins have to be equal to these unobserved consumer types.

So again, the way this works is they only have these three parameters here. This is a method of moments problem. And for each time they evaluate the objective function, they directly impose this correspondence between observed margins and unobserved consumer types.

So what do they get out of this in the end? So one thing that's typically, of course, when people compute structural models, they use them to compute counterfactuals. Here, the counterfactual is already observed in the data. So the additional thing that you get out of computing this model is that you can learn something about the unobserved search costs of consumers. They estimate that this is about $1,700 for an average loan size of about $152,000.

And because they also get lenders' costs from this, they can compute their markups, which are at the order of 5% to 10% And so, again, this model, I mean, it's constructed in this way, but it directly maps to this empirical observation. The consumers at the bottom end of the price distribution are hurt more by the merger than consumers who are searching much less.

I think this is a very nice paper, because again, it uses an actual change in market structure that's observed in the data. So an actual merger. It highlights what I think is a new angle on the effects of mergers, is that different consumers fare very differently. And that could potentially mean that we underestimate the market power effect of a merger.

Because if, let's say, consumers don't search very much or we have markets in which search costs are very high, and we do, let's say, a simple before-after merger analysis, we might observe that prices don't change all that much. But that could be because consumers are just not searching very much. And so to the extent that we want to learn what's the market power effect of a merger, it's important to account for that.

And it provides this what I think is very tight link between the treatment effect analysis and the structural model. And that, of course, required some, as I think I've pointed out, fairly strong assumptions. In particular, to come back to those, the fact that there are no information asymmetries between the borrowers and the lenders means that search is actually not happening in equilibrium.

So the lenders-- think about the quote of the initial lender. The lender is always paying consumers such that they do not find it worth their while to go out and get additional quotes. So that means if you want to, for instance, quantify the, let's say, lost economic surplus due to search costs, this wouldn't be the right model for that.

And we might think that in some markets, the fact that some consumers have lower search costs than others can have disciplined prices in the market overall. So you had seen this effect, for instance, in the style model that if you get more sophisticated consumers in the market, we move towards more competitive pricing. And this kind of effect you don't get here because of this perfect price discrimination where every type gets kind of his price.

So I think that's a potential shortcoming of this model that is kind of interesting to explore. But I think they had good reasons to construct it this way, because it provided this tight link to the empirical specification. Any questions or comments about this paper? Yes.

STUDENT: I guess here the search costs are just based on each consumer. There's nothing that the firm does to change the search costs. But it seems like in a mortgage market, I could imagine that maybe the firm does less advertising or I make you walk into my office rather than call me. And that would be, like, if [INAUDIBLE] strategy for changing the search cost depending on how much competition there is. And so what is the reasoning for trying to think of firm strategies just in the prices that they give consumers rather than any way in which they could be changing the search costs?

TOBIAS SALZ: So first of all, there are papers mostly in the marketing literature that actually do what you suggest. They think about advertising as basically firms providing a search, a free quote, to consumers. There's actually a fairly recent paper by Andrea, Siminoff, and Raluca Ursu, and some others that do exactly that, where they try to see how does a free quote or free presentation of a product influence a consumer's search strategy?

And they do this with online data. So people have thought about this. And I think it's interesting. I think in part, it's tractability. But yeah, that's a good thought. Any other questions about this?

So I'll probably not get through all of this, but I just want to very briefly talk about my job market paper, which ties in quite nicely with this previous paper that we discussed, because in some sense, one of the kind of economic differences that I'm after here is precisely getting at these spillover effects of having more sophisticated consumers in the market.

So again, the idea, which is, for instance, in the Salop-Stiglitz model also, that when there are sophisticated consumers in the market, both sophisticated and unsophisticated ones, and firms can actually not tell them apart. Then the fact that some consumers are searching disciplines the firm's pricing, because they don't know which kinds of consumers they're dealing with. And this effect was kind of completely shut down in the previous paper.

And so in particular, I want to look at this effect with an application to intermediaries. So intermediaries are in this setting firms that help buyers search. And if you think about who are the buyers that would hire somebody else to conduct their search for them, it's going to be buyers that have very, very high search costs.

So you could construct a model where there's some sorting very closely related to the Hong and Shum model that I showed you, where you have basically these different consumer bins. Consumers with higher search costs take more draws from the price distribution. Sorry, take fewer draws from the price distribution. Once your search cost becomes too high, you instead use an intermediary.

And so what we get out of this is that the highest search cost types are kind of leaving the market. They're going instead through the intermediary. And firms in their pricing will kind of condition on the fact that in equilibrium, they're not encountering the very highest search cost types. So again, think about the Hong and Shum search-cost distribution. You have at the very top the people that leave the market. And that could lead to lower prices for all consumers in the market.

And so this, of course, only works if you now do have this private information, if firms do not know whom they're dealing with. And that introduces some technical challenges in how to deal with that. And so basically what I do in this paper is construct a model that essentially puts together the Hong and Shum search strategy of consumers, where consumers have heterogeneous search costs and decide how many draws to take from the price distribution, with a model of suppliers that also have heterogeneous costs for providing services. And in these kinds of settings, you now have unobserved source of variation on both sides of the market.

So again, one contribution is here to construct a model for these kinds of decentralized markets where firms don't know what type of consumer they're dealing with. And one data contribution is to find a setting where you actually find prices for decentralized markets. So typically good settings are those that are regulated. So privacy, in the case of the mortgage paper. The reason that they had this data is because people needed to get insurance. So they collected what is essentially administrative data, because the lenders had to submit this to a regulator.

And so a similar thing in my setting here where I'm going to use data from the New York City trade waste market, which were very idiosyncratic reasons, which is basically organized crime in the city of new York was starting to be regulated. And all the firms have to submit their data to a central regulator. And so you can actually see individual level contract data that allow you to make some progress on these kinds of markets.

OK, so this is the non-residential waste disposal market. So trade waste market is a fancy name for that. And so what is really striking about this market is that there are a lot of sellers in New York City relative to most other cities in the US. So in most other cities, you have large companies like waste management and others that are supplying the entire city. Here you have about 100 sellers in some zip codes of Manhattan. You can call up to 35 carters. So there seems to be this very fragmented and, from an outside observers point of view, redundant supply.

The market is regulated by the Business Integrity Commission, which was established by the city in 1995. And again, was a response to the property and racketeering system that has been in place in New York. Famously, many crime shows and movies, that the waste market was one of the predominant industries where organized crime was active. And that led the city to actually collect all this data.

So I want to use my last few minutes to just show you quickly the model and how it works just so you get a sense. It's a bit fast, but just to get you the idea of what is different about this model. So here again, we have buyers. These are businesses in New York City that need to contract for their waste disposal services. We have sellers. These are carting companies. And then we have brokers in the middle. And some buyers could decide to delegate to a broker and the broker instead searches for them, but by searching here means it's an auction.

And so what is challenging about this model is that we have basically two sources of unobserved variation that gives rise to prices, which are unobserved costs of businesses to provide service and the unobserved search costs of buyers to search, which are here called kappa.

And so if you, for instance, contrast this with the Hong and Shum model, in Hong and Shum, the cost of sellers was assumed to be constant. And the heterogeneity that you observe in prices basically maps to an unobserved distribution of search costs. So here you have to disentangle the two.

So let me skip some of this. I told you about this. I just want to briefly show you how the model works. So we have buyers that have heterogeneous search expenses, and they can, like in Hong and Shum, decide how many draws to take from the price distribution or whether or not they're using a broker instead. And so because you know the Hong and Shum model, I can skip ahead a bit here. So you get, again, a partition of the search cost distribution, but those with the very highest cost, they're going to use a broker instead.

So as before, we want to keep track of buyers that are the mass of buyers between these marginal types. This gives rise to these search weights. So these are the buyers that search m times from the price distribution. And so because the seller's cost is also heterogeneous, we don't get a mixed strategy pricing equilibrium, but we instead get for any different cost that a seller could draw, we get a bidding function that maps the seller's unobserved cost to a price.

And so just to show you this briefly, I know I'm going fast here, but we now get a profit maximization problem that looks like this. The sellers are maximizing for any given consumer that contacts them their margin times the weighted average across different consumers that they could encounter would take n draws from the price distribution times the probability that they're selling to that consumer.

Now, what is different here is that these marginal costs are not uniform as before as in Hong and Shum, but they're drawn from a distribution. And so what's nice about thinking about this problem is that this actually leads to a closed form solution for the bidding function that is like in a first price procurement auction.

I know we'll do auctions later in this semester, but this is a first price procurement auction solution to an IPV problem. And this is essentially just a weighted average across what you would bid if the number of competitors were known. And so here the number of competitors depends on the consumer search strategy.

So just to wrap up and put the pieces together, so basically this kind of takes the Hong and Shum model and then introduces heterogeneity on the supply side and kind of maps this essentially back to what is an auction problem. And I'll leave it here. I'll probably tell those who come to [14.]273 some more about this. But yeah, that's it for today. Thank you.