[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK. Welcome back. So if you remember, last class, I was talking about search theory, and discussed two main things. One is that search costs can complement product differentiation and raise prices in markets. And the other is that in Bertrand-like markets, search costs can also lead to price dispersion.

In most of this class, we've had a pattern of two lectures, one where I do the theory, and then one where I talk about the empirics. On search, it's a big topic. We decided to do three. So I did the theory last time. Today I'm going to talk about reduced form evidence on empirical evidence on search, and then we're going to have a third lecture on search, where Tobias is going to talk about the more structural work. And that will come later.

So anyway, empirical work on search-- again, why do we do empirical work in IO? We're often generally trying to do two different things, or maybe a combination of two things. One is we're trying to get evidence on the theories, whether they're applicable, whether they're what might be correct or incorrect, or all theories are correct. But what is it about the theories we might want to modify to have it fit applications? And the other thing we're trying to do is just take the theories and apply them to particular empirical examples that we care about and estimate parameters of interest, how big or small effects are.

So I'm going to talk about four papers today. The first of these is of the just gaining insights in the theory literature. Generally, when we're trying to gain insights in the theory, we're often looking at unusual situations because we want to look at markets that are particularly simple, where the theory is particularly stark to try to understand its predictions, whether they apply.

Particularly, Alan in his paper is talking about-- trying to understand, do we see the kinds of price dispersion predicted by Stahl-like models of heterogeneous consumer search? And is it true that frequent violations we see of the quote, "law of one price" are because of consumer search costs leading to heterogeneity?

And so in this paper, what Alan was trying to do is think of an application where he could argue that the Bertrand assumptions fit very, very well and examine in a model where Bertrand assumptions fit very, very well, do we see price dispersion? And can we provide evidence that price dispersion seems to be due to consumer search costs?

You'll see that this is a paper written when Alan was a PhD student at MIT. Very much it's a paper that has sort of minimal data, self-collected that anyone could write. So his idea was that we could get evidence-- a good example to think about Bertrand-like competition is prescription drugs. One of the nice things about prescription drugs is that a bottle of amoxicillin tablets bought from pharmacy A and a bottle of amoxicillin tablets bought from pharmacy B should be completely identical products. People should not-- people should know that these things are exactly the same product, that we're just buying it from different stores. So we could think of people as having Bertrand-like preferences.

One of the other nice things about prescription drugs is you're always struggling in empirical papers. How do I get the number of degrees of freedom to actually estimate something? The great thing about prescription drugs is there are hundreds of different prescription drugs, hundreds of different prescription drugs given hundreds of different observations.

So what Alan did for this paper is if you think about the Northeastern United States, so we're in Massachusetts, which looks something like this. You've got Connecticut here, Rhode Island, New York State. Alan noticed that New York in the late 1990s had a law that pharmacies had to post all of their retail prices for the 200 most popular prescription drugs on a poster that was in displayed in large format in the stores. Therefore, what one could do is go to a pharmacy in New York and very quickly find out what prices they were charging for every different drug.

In the 1990s, people did not have phones in their pockets with high resolution cameras. So you could not go there and just take a picture of the price photos. So what Alan did was go to New York with a clipboard and just copy down off these price posters what prices the pharmacies were charging for all the drugs. And in particular, he does this in two small towns, Middletown and Newburgh, New York. And these happen to be located where if you drive from Boston along the Mass Pike, and then get on 84, and then go to New York, these are two of the first towns you come to in New York State.

One thing that was nice about them is he wanted to get all the prices in a market. Middletown and Newburgh are both small towns. They have 10 or 12 pharmacies in them, and then there are no other pharmacies within a 15-minute drive or 20-minute drive away. And so you could think of anyone who lives in Middletown as shopping-- having the options to be buying prescription drugs all the pharmacies in Middletown, and in Newburgh, they can buy the things in Newburgh. And I think I've said all these things.

Oh, I guess another thing about how the world was different-- another thing about how the world was different in late 1990s is in the late 1990s, most people did not have prescription drug insurance. And so at the time, the prices that the pharmacies charged for a prescription drug was what people-- customers would be paying. Customers would have the same incentive to shop for drugs they would have for any other product. Obviously, today, many more people have prescription drug insurance. The prices at pharmacies would not be what they're paying. But at the time, this thing had this advantage.

So anyway, here's an example of the study design. So you can see that-- this is not the actual poster, but there was a poster that was very much like this. I think it was roughly 150 drugs, but then many of them had generic as well as branded versions. And so you were getting over 200 prescription drugs.

And so there would be this big poster. It would list Accupril, 10 milligram tablets, Accupril, 20 milligram tablets, and so on. And then it would have the prices there. And so Alan could just go into the pharmacies, copy down the prices listed on this poster, and get an example of what they cost.

And different drugs have different characteristics. And so what he's going to be able to do is not only say how much dispersion is there within a drug, but then how much dispersion is there across drugs? Sorry-- how does the dispersion vary across drugs with drug characteristics?

One thing that's also nice about this application is you could say one reason why drug prices should be different across pharmacies is because their drugs are bundled with pharmacy amenities. Some pharmacy has better customer service, is cleaner, and might be quicker to get in and out of. People would pay extra to go to that pharmacy. But the idea is having 200 different drugs, you can put in pharmacy fixed effects, and say, the pharmacy may be nicer, but there's not something like, this pharmacy is a nicer place to get amoxicillin, and that pharmacy is a nicer place to get cephalexin. Pharmacy fixed effects ought to account for most of the amenities, and lets him control for that.

Not much sophistication in the analysis, so I won't go through how it's really done. But let me just talk about some of the basic facts. So first basic fact, there's substantial price dispersion. The difference between the highest and lowest price for drug across the drugstores in a town is $13 for the average drug. Some drugs have less price dispersion. 10th percentile range is $5. 95th percentile range is $25. So substantial price dispersion, obviously, it's not enormous, but $13 is about what you can think of.

The second finding is that the price dispersion doesn't appear to be due to pharmacy fixed effects. He argues this in a couple of ways. First is just this simple table where he puts for every pharmacy in these towns, how often is their price in the bottom third of the distribution? Here in Middletown, there are 10. So how often is their price one of the three lowest? How often is it in the middle? How often is it in the three highest?

And for many of the pharmacies, like the ones I've highlighted, Eckerd, Immediate, and K-Mart, there's a substantial number of drugs where they're one of the lowest priced pharmacies, a substantial number of drugs for one of the highest priced pharmacies. There are a couple exceptions. Rite Aid seems to be a terrible place to buy drugs if you live in Middletown. It's almost always one of the highest three prices-- also, almost one of the highest prices in Newburgh. Price Chopper and Walmart seem to have lower prices than the others.

But if you think about how important are the pharmacy fixed effects, he has a simple table where you regress the prices of drugs in a pharmacy on drug fixed effects. That has an R squared of 0.907. You include drug and pharmacy effects as an R squared of 0.938. So what that's saying is the pharmacy fixed effects or the amenities are accounting for about 1/3 of the price dispersion, which leaves 2/3 of the remaining price dispersion. So 2/3 of the $13 appears to be pure, idiosyncratic variation across pharmacies, not correlated with pharmacy fixed effects.

And then Alan also argues in the paper even that may be an underestimate of how much of it is pharmacy fixed effects, because he looks at the expensive pharmacies, like goes to the Rite Aids, and goes to the Price Chopper. And his informal evidence of, how friendly is the pharmacist? How clean is the place? How long is the line? He would rather go to Price Chopper than go to Rite Aid. So in some sense, this high price at Rite Aid may just be part also of the idiosyncratic price dispersion, not part of the bundling with amenities.

The main thing that Alan does to then argue that this price dispersion is related to search costs is to sort of-- you want to think about-- well, let's go back to the theory. If you remember, what did I teach about price dispersion? If you have a market and you have-- if we have very low search costs, and this is price equals cost, what you would see is with very low search costs, we get a distribution that's fairly tightly distributed around-- we get a distribution of prices that's fairly tightly distributed around costs, because if anyone sees a high price, they're going to be tempted to go back and search again.

That gives us a search cost distribution that's low. Prices are close to cost, and then not very dispersed, because if prices were very dispersed, anyone seeing a high price would just go search again and try to get a low price. So we get very little price dispersion.

When we have high search costs, what you get instead is more dispersed distribution that may look something like this, where a few firms set low prices. Many more firms set high prices. Even the consumers who see the high prices don't want to search again.

So we get two different things. We get a higher average price, and we get more dispersion, whereas here we get a lower average price and less dispersion. So that's what Alan would like to say, is, can we observe drugs that have lower search costs versus higher search costs and find out do the ones that have lower search costs have lower and less dispersed prices?

It's hard to think-- you can think of reasons why some drugs may have higher or lower search costs. So for instance, drugs that treat conditions that are acute and very painful or make it difficult for you to shop could have high search costs. Drugs that treat sexually transmitted diseases might have high search costs because you're embarrassed about the condition. You don't want to go tell every pharmacy in the town, call them up, and say, what would your cost be for buying this prescription?

But what Alan comes up with is this idea of frequently purchased drugs. So you have some drugs, like antibiotics, that you're typically buying and you're going to use once. You have other drugs, like high blood pressure medications. High blood pressure medications, almost everyone who's taking them is going to take them-- buy them every month for the rest of their lives.

And so if you think about shopping for an antibiotic versus shopping for high blood pressure medication, it pays to shop around for the high blood pressure medication because you're going to be able to amortize those search costs over the rest of your life. Whereas the antibiotics, you're just going to buy this thing once. You're never going to buy it again. The benefit of search is smaller.

So what Alan does is think about frequently purchased drugs as drugs that have effectively lower search costs and examine whether the range in prices and the level of prices are lower for drugs that are purchased more frequently. The range is what's done here. So he has four different models-- four different measures of the range of prices, the range, standard deviation, the residual range after controlling for some pharmacy fixed effects and other things, residual standard deviation. In all four cases, he shows that the more frequently purchased drugs do have less price dispersion than the frequently purchased drugs.

He also chose to have a table for this. I do not. He also shows that the average markups for these drugs relative to published wholesale prices are lower for the frequently purchased drugs. And then a third fact, as I said, we expect for idiosyncratic reasons, some drugs to have lower or higher search costs apart from this sort of purchase frequency. And what he shows is that drugs that have unexpectedly high average markups also have unexpectedly high dispersion, suggesting that there's also unobserved variation in search costs, and that unobserved variation in search costs is driving some of the price dispersion and price levels.

Anyway, so I think it's a very nice paper. It's providing very simple evidence saying that these theories we have, like Stahl's model of price dispersion, do appear to be applicable, and do appear that the price dispersion they predict is there. And it appears that the price dispersion is search cost related.

Stango and Zinman are on the opposite side of the primary second motivation, trying to look at a question that we really care about and understand. Does-- do models of price dispersion help us understand what's going on? And how big are the magnitudes of price dispersion in this application?

The particular thing they're talking about is credit card debt. Credit card debt is something that many people feel is unfortunate. There are many people in the United States who pay much more on credit card interest and have gotten themselves into substantial trouble.

There was a set of papers in the 1990s set off by this *AER* paper by Larry Ausubel raising what he called the credit card puzzle. And the question is, why is it that credit card interest rates are so high, given that there are a million banks in the United States? Why is it that this competition between the huge number of banks does not dissipate profits in the credit card industry? And how do we get profits that are so high for the firms?

The other thing that Ausubel raised was actually non-dispersion intuition, which is that Ausubel reported that not only are credit card rates very high, but every big bank charges very high rates on their credit cards. So there's not price dispersion. He would have thought-- he'd sort of say, why is it that all these banks are charging very high interest rates? Why doesn't one of them undercut the others and try to compete with them, bring interest rates down?

Ausubel also reported that rates were fairly insensitive to interest rates. So his data covers times when-- there are times when interest rates on mortgage loans are 10%, times when those rates are 2%. When interest rates on Treasury bills and mortgages are changing, credit card rates seem to be fairly insensitive. And there was another puzzle he raised is, why is it these rates are not-- why is it they're so fixed, and there's so little dispersion? They don't vary over time or across banks.

So Stango and Zinman are actually going to argue that the basic facts that Ausubel had us think about were wrong. It's not the case-- it is true that rates across banks tend to be fairly similar, but what's not true is that rates across individuals are true. And so there's this hidden thing that Ausubel was missing, which was that all of these banks have high average interest rates, but all of these banks are also offering many different interest rates to different clients. Some of their clients are paying low interest rates, some of their clients are paying high interest rates.

And so what Stango and Zinman are trying to think about is document that there's substantial price dispersion, and think about search costs as a potential explanation for why it is that we see these high interest rates. So anyway, some basic facts in their paper, and I should say, where does their data come from? So this is not the sort of amateur data collection we see in Sorensen's paper. This is the data collection where you make connections with some firm that has access to very high quality data.

So in particular, there's a company called Lightspeed. And there are a few companies like this. What Lightspeed does is it has a panel of consumers who basically open their lives to Lightspeed in exchange for some type of payment. So in particular, there are 4,312 customers in their data.

What these customers do is they give Lightspeed access to their credit card statements. So every month, Lightspeed gets their complete credit card statement, everything they bought, how much, what they paid, how much interest they paid. Did they pay a late fee? It's basically like they've got the credit card statement for 4,312 people.

They also have information on the people in their panel. So when people join their panel, they complete a survey that gives them a lot of demographic information. They also know their consumers' credit scores. They've got this from-- I don't know-- Experian or some other credit reporting agency. So they have some measure of the credit worthiness of people, which you would think would be a driver of the heterogeneity in interest rates being paid.

So let me start from some basic facts. The first fact is there's tremendous heterogeneity in interest rates across consumers. If you look at the interquartile range, just between the 75th percentile and the 25th percentile, it's 800 basis points, or 8 percentage points per year. This is omitting anyone who has teaser rates and who pays in full. So these are people who are paying interest on their credit cards every month. There's just tremendous heterogeneity there.

Of course, there's going to be heterogeneity in credit card rates, because some people are bad credit risks. Some people are good credit risks. You would think if you were running a credit card company, you would give lower interest rates to people who are going to pay you back, a higher interest rate to the people who you think might default.

What they find is that default risk explains about 40% of the interest rate variation. Other factors, like offsetting rewards, demographics, explain very, very little. So roughly 60% of this gap is due to what appears to be pure, idiosyncratic heterogeneity. People who are similarly situated, equally likely to pay the bank back, some are paying much more than others.

And a third fact I'll talk about later is that there's also substantial within consumer variation offers received. They have a separate second data set where they have people open-- they know all of the credit card mailers that are going to particular individuals. People in their data-- I don't know if this applies to you-- get multiple credit card offers per month. They look within one month what credit card offers people get. And there's substantial variation in the offers if you open them up.

So let me just sort of-- some summary statistics, and this is a case where I think part of the interest in this paper is this Lightspeed data gives us this view of what is it that consumers are doing. And you have sort of shocking statistics about problems people have.

So they divide the data into the quartiles of the average balance that people have on their credit card. You were probably told that it's good practice to pay off your credit card bills every month, and simply not pay any interest. And what they find is that the 25th percentile balance is $499, which means the fraction of people in their data who pay off their bill every month is less than 25% So it's close to 25%, but most people have a balance on their credit cards, and run a balance on their credit cards every month and never pay it off.

What's the upper quartile? The upper quartile is people with credit card balances between $4,586 and $62,000. People in this quartile, on average, have a revolving balance of $11,000 on their credit cards. They're paying $2,000 in interest. And obviously, these are the people that people really worry about and people who are making big financial mistakes. And is there something we can do to help people avoid getting stuck in this situation where you're spending $2,000 a month on interest?

You may also notice that the people who are spending $2,000 a month in interest, what's their income? A quarter of them have incomes below $45,000. So being income below $45,000, paying $2,000 a month on your credit card, these are people who've made very big mistakes. Some of them have incomes of above $100,000, but it's still substantial.

Also, noteworthy is in some sense, some of the U-shaped relationships here. It's actually people in the bottom quartile and people in the top quartile who have the best credit scores and who are charging the most every month. Some people in the second quartile-- people in second quartile are spending less per month, have lower credit reports, lower credit scores. These people are probably constrained in how much they can borrow. But it's the people at the top, actually, have fairly good credit scores, and the credit card companies enjoy being able to charge them $2,000 a month in interest.

So this fact I told you about, there's tremendous dispersion in credit card interest rates. Here's the basic raw information on what people are paying. And you can see it's not based on your revolving balance. So people in the lowest quartile, there are people paying between 12%-- there are people paying 12%, and people paying 26%.

So this is the 10th percentile. This is the 90th percentile. People in the upper quartile are people paying 11%, people paying 26%. It seems to be very similar across those four groups. For every level of revolving balance, there are people paying very low interest rates, people paying very, very high interest rates.

As I said, the second fact they have is, are people paying very different interest rates because they have very different credit scores? We see here is the blue dashed line is the distribution of interest rates people are being paid relative to the average rate. So 0 is here. So there are a number of people paying interest rates 5% below average. There are a number of people paying 10% below average. Also, many people paying 0 to 10% above average. Not many people pay more than 10% or 12% above average. But the blue is a distribution of interest rates people are paying relative to the average.

The red is what you get if you say, now I'm going to control for your credit score and ask, once I've controlled for your credit score, what are you paying relative to the average for someone with your credit score and other attributes? And what they find is that this explains some of the variation. The red distribution is more tightly distributed than the blue distribution, but it's explaining perhaps 40% of the variation. So some of these tails go away, but we still have very substantial differences in interest rates paid by people who appear to be similar on all attributes.

So then the next thing the paper does is try to say, can we tie this to search intensity? And can we find that people with lower search costs-- the people who are more likely to search are the ones paying interest rates dowr here, and the people who are less-- have higher search costs end up getting stuck and paying high interest rates up there?

It's not as clean as Alan's paper, where Alan has this exogenous characteristic of a drug that shifts its search costs. But what they have as a measure of search costs is just a self answer to a survey question. They ask consumers-- so they not only have access to the Lightspeed panel data. They have the ability to ask questions of consumers in the Lightspeed panel data. And they ask those consumers, how likely are you to look at a credit card offer you get in the mail?

And you might think that that would be a measure of search costs. People who have low search costs, people who like opening credit card envelopes, and look, thinking about them are people with low search costs. People who dislike opening their credit card offers they get in the mail are people with high search costs. Is that going to explain things?

When you think about that, there's this obvious endogeneity concern. You ask people, how likely are you to open a credit card offer you get in the mail? Well, if you know you have a bad credit card and are paying a high interest rate, you may be more likely to say, yes, I'm going to open up things, because you're looking for a new credit card.

Whereas if you already have a really good credit card, you may say, no, I don't open them up. And you're not opening them up, not because you have high search costs, but because you know you already have a very good deal.

So what they do is-- the first thing they do is they just run this OLS regression. And they regress the APR that you end up paying on your self-reported search intensity. And we get a small negative coefficient, saying, people who search more actively are paying slightly lower rates, but then they sort of-- it has that endogeneity, so they run this IV regression.

The IV regression says that we can think about gender and marital status as instruments for search intensity because the offers that you get in the mail should not differ across gender, across marital status once you control for credit score and other factors, because it's illegal to discriminate against women or against married or unmarried people. Therefore, those things should relate to your search intensity, but be unrelated to the set of raw set of offers that you get in the mail for credit cards.

When they put that IV in, now they get a very large negative coefficient, saying, people who search more intensively are paying lower interest rates. I mean, it's a very large gap between the OLS and the IV estimates. And then the standard error really blows up relative to this.

So you wonder-_ again, it's only a slightly significant result. Why is that? Is the reverse causality really so large that it causes this? But anyway, that's the evidence they provide, that it does seem to be the dispersion that people are paying is related to some characteristic that does their search intensity.

So you'll see at the end of this lecture, I'm going to come back to another paper that further examines this question, but I think it does clearly reframe the debate relative to what you thought from Ausubel and other early papers to say that there is just substantial price dispersion in credit cards. Search costs seem to be a part of why it is that it's so hard for firms to compete in this market is if you send out lots of credit card offers and people don't open them, people don't consider buying them, that may make it one of the big factors that's obstacles to eliminating the high markups that we see.

So moving on, I'm going to spend much longer talking about a paper of mine. I talk about it in part because I know it better. It's one of my favorite of my papers, or papers, mine with Sara. This is a paper where we're both trying to-- we are trying to-- it's a more talking about the theory paper where we are trying to understand what causes price dispersions and markups.

I see this-- we saw it at the time as a paper that was looking forward. The question was that in the early 2000s, search engines, price search engines had recently been invented. There were a few markets, like airplane flights, hotels, rental cars, where many people use price search engines. Most other markets, people don't still use price search engines.

But our question was, what's going to happen in the world if price search engines become much better? First of all, how are the airplane flights, and the hotel, and rental car markets going to adjust when flights get-- when search technologies get better? And what's going to happen to the rest of the economy when search technologies get better?

I mean, I think somehow this still seems topical 20 years later, when we see these amazing AI advances and chat robots, and whatever. You wonder, am I going be able to-- are we going to get to a point where I can just type into Google, what's the best place for me to buy amoxicillin? And have amox-- and have Google say, you're living-- I can see that you live in Newton, Massachusetts. For you to buy amoxicillin, the best pharmacy to go to is this one that's on your drive home, and it costs this much there, and this much somewhere else.

Or are we going to get to the point where I can say, I want to buy a tennis racket, and then Google just tells me, given what I know about you, this would be the best tennis racket for you to buy. And this store has it at this price. This is where you should buy it, and just get it shipped directly to your house. So who knows whether we will ever get to that point. But the question, I think, still remains, how do we expect markets to change when search costs-- when internet technologies improve. Search costs go down. Are we going to see a collapse in price dispersion, where formerly high dispersed prices become very low tight prices?

And we mentioned in the paper the potential Bertrand paradox. Real retailers have substantial fixed costs. The fixed costs of running a store are something like 20%. So if you can't earn 20% markups over your wholesale costs on goods, you're going to go out of business. Even if you're Walmart, you're super efficient, you need to have 20% markups of your wholesale cost to cover all the fixed costs of having a physical store, and things like that.

And so the question is, how is retail and other business going to survive when the internet makes price search more efficient? And I guess, one of the answers we're going to say in this paper is that, models of search traditionally thought of consumer search as a one-sided problem, where these consumers have these search costs. The search costs are fixed. And then they search, and that fixed search cost determines the markups.

The thing we're going to think about is actually two-sided. You have the search engine is trying to make search more efficient, and the consumers are trying to make their search more efficient. And at the same time, the stores are on the other side, trying to make search less efficient. And so we can think of search costs as not being exogenous, or not being just a result of technology, but being an outcome of a two-side-- equilibrium of a two-sided game, where some people are trying to make search more efficient, other people are trying to make search less efficient, and this balance between search and what we call obfuscation ends up determining price levels and markets.

So as I said, this is a paper where we're focused on this theory of, how would obfuscation work? We go to a particular narrow corner of the internet, where everything is much simpler, and it makes it much simpler to do the analysis than it would if we look at a more popular product. And so in particular, we look at a price search engine called pricewatch.com. I think it still exists, although it's long past its heyday.

But in the early 2000s, Pricewatch was this simple database search engine that people who were sophisticated would use to find things that they needed. So for instance, if you were to go to the computer guys in the MIT Economics department and say, I have a computer. I'm getting low memory errors. Could I just buy some extra memory and plug it into the-- plug some extra memory modules into this part of my circuit board and increase the memory capacity of my computer? People who worked in IT offices might use pricewatch.com to find an inexpensive place to buy memory that you could add to your computer.

What did it look like? And again, this is-- think 2000. The world was a much more primitive place. People had lower bandwidth on their home internet. Things just had simple text-based designs. So this was Pricewatch-- very proud that it was established in 1995. And you can see the pixelated graphics there.

But anyway, you'd go to pricewatch.com. And it would say, these are the things you can buy through our site. You can buy computers, PCs with or without an operating system. You can buy CPUs. You can buy motherboards. We're going to be buying memory. So RAM memory versus flash cards, we're going to be buying RAM.

So anyway, I go to this page. I click on RAM. When I click on RAM, it brings up this page, which shows me all the different kinds of RAM I can buy. Again, you need to be somewhat sophisticated. Do you need DDR memory? Do you need notebook DDR2 memory?

These are all the memory modules you can buy. They're going to have various descriptions. So the 512 megabytes, 4 gigabytes, 2 gigabytes, that's the storage capacity of the memory module. This 4000, 3500, 3200, those are describing the speed with which the memory communicates with the motherboard. The DDR, DDR2, those are the other aspect of the thing.

Basically, you need to know what type of memory you need for your computer. My computer needs DDR PC3200 memory. It is a choice then. Do I buy 4 gigabytes, 2 gigabytes, 1 gigabyte, based on how much I want to spend, how much memory I want to add? But the site is designed for people who are sophisticated in knowing, what memory do I need for my motherboard?

This is now getting to actual data. So this was taken a bit later, when you can see people buying 4 megabyte and 2 megabyte-- or gigabyte modules. Back in 1999, what people would buy is either 128 megabytes of memory or 256 megabytes of memory.

So I clicked on one particular type of memory, 128 megabyte modules, that follow the-- PC100 would be the-- 100 would be the speed which it communicates with your computer. So you know you need PC100 memory. You want to buy 128 megabytes.

You click on this link. And it brings up this page. This page just has a set of offers for where you could buy the memory from. I can buy it from Computer Craft, Inc, from Connect Computers, from First Choice Memory. This retailer is located in Florida. This one is in California, and so on.

And then the price at which they'll sell me the memory is 68, 69, 70, 72, 74, 74, 74, 75. And then how much they'll charge me for shipping, it's usually-- at the time, Pricewatch had a cap. You couldn't charge more than $11. And so many of the firms are some number that's slightly less than $11.

Again, it was a primitive period, where if I wanted to buy the memory, I could then click on Computer Craft Inc.'s website. It would take me not to-- it's not like this would add the memory to my cart. I would click on that button. I would typically get to the Computer Craft site. And then I'd have to search around on the Computer Craft site to find where is their 128 megabyte PC100 memory. Can I find it at that price of $68 that Pricewatch is telling me that they offer it at.

Pricewatch was not scraping these numbers. The retailers were having to enter these numbers into a database at Pricewatch. So if you're one of these retailers, you would get up in the morning, decide what your price is, upload the new price to Pricewatch, change it on your website, and that's what you would charge.

So basic observations on Pricewatch is there was a tremendous amount of competition on Pricewatch. Prices were very low. At this time when this memory module cost $68, if I'd gone to Best Buy to buy the same thing, I might have paid $100. So markups were very low on Pricewatch.

But there was some dispersion. And it was clearly not at all the frictionless ideal of, I click there. Everyone's in Bertrand competition. I get exactly what I want. Two experiences were common when you used it.

First, as I said, you didn't go directly to the landing page. Even when you did go directly to the page where you clicked on, you would find these pages that would be annoying and time consuming to buy what you wanted to buy. So this was an example from a bit later, where I clicked on a offer from topshop.com to buy some memory module that was supposed to cost $53.81. I clicked on a link, saying, buy this module for $53.81. And they give me this page that tells me-- and there's actually more than this page. This page saying, price with all the selected options is $90.36.

And so, why is it that it's $90.36? What they've done is this page, it's got all these other things that you can click on to add to the cost of your module. And some of them are just pure annoyances. Somehow, they've pre-checked for me "Bonus buy, 10 pack of hand thumbscrews for my case."

So if I had hand thumbscrews to open up the computer case in which I wanted to change my memory, why I would need new thumbscrews, I can't imagine. But anyway, for $4.95, I can get those 10 thumbscrews. I just have to go through, read that, and say, no, I don't want that. Uncheck that box, and it takes $4.95 off the price.

Others of these are things that it's less clear that I should uncheck. So this is-- take advantage of these special offers-- memory upgrade, CAS 2.5 upgrade. Improve performance and help things. They're charging me $6.35 from that.

This is another example-- six layer stability, more layer, better design-- not in perfect English, $8.37. But then they also have these others that are more concerning. Pre-test, standard pre-test, avoid costly return agreements, $6.97. So they're telling me is that if I want them to test the product-- they're shipping it to me. Before they ship it to me and make sure it works, I have to pay $7 to get them to test the product they're going to sell me, and make sure that it works. And if I don't ask them to test it for me and make sure that it works, they're going to charge me to ship it back to them.

And so it may be that you look at this offer and realize this offer is not what I thought this offer was. I should think of this offer as being $7 higher than this offer was. And in some sense, I'm going to have to just spend a long time reading the fine print of Tough Shop's offers, where I trust, do I want to buy this product? Do I not want to buy this? And it's just going to take me a while to shop.

The second thing that was common was offers that were not designed to just delay and annoy you, but ones where it was clearly aimed to get you to shift away from buying something you bought into something else. If you remember back to when I talked about competitive price dispersion, I talked about add-on pricing. That is, if you have two firms selling two products, models of price dispersion were totally well-- sorry-- models of price discrimination were totally well with differentiated products and multiple firms.

And you can have two firms competing against each other. They advertise all their prices. They sell low quality goods at low prices, high quality goods at high prices. And that can be rational if the consumers for the low quality goods are more price sensitive across firms. And so you have an equilibrium where it's like the competition on two-line model. For low quality goods, firms are very undifferentiated. For the high quality goods, they're differentiated. These goods are sold at a low markup. These goods are sold at a high markup.

And one thing I said when I talked about that is one way in which firms in this model can raise prices is you can think of another game where you advertise your low quality product, and then you have an unadvertised price for a higher quality product. In that game where you only show people the price for the high quality product once they show up at the store, I gave this model that said, markups are going to be even higher in equilibrium.

And the reason is that firms are going to charge the ex post monopoly price for the upgrades. That causes a wedge between-- when you force the prices of the low and high quality goods further apart, there's this adverse selection effect, where you don't want to get-- you don't want to sell to the cheapskates who are going to buy your rental car for $19.95 a day, and not get the insurance, and not get the car seat.

And because there's this adverse selection effect where it's the cheapskates who don't buy the insurance and the rental-- and the car seats, you then want to dump-- you go from wanting to undercut the rival to get consumers to wanting to overcut the rival to dump the cheapskates on them. Markups for both products can end up-- markups for the sort of-- average markups can go up in equilibrium.

We see what looks like an add-on strategy here. So you've clicked on this product I want to buy. I clicked on, I want to buy OEM 512 megabyte memory. They actually put the check mark in the correct box. This is what I said I wanted to buy. But then they tell me that I can spend $15 more and get this product, or I can spend $25 more and get this product, which I'm going to think of as add-ons.

What do they do? They tell me a lot about why the better products are better. So for instance, this one is CAS 3 latency. This is CAS 2.5 latency. This is CAS 2.5 latency. This one has a four-layer board. These each have six-layer boards. These have OEM DRAM downgrade chips, whereas these have industry standard chips, whereas these have even better brand name chips. Samsung and Micron are major brands.

These have a restocking fee. These have no restocking fee-- satisfaction and compatibility fully guaranteed. In fact, one thing that-- for fun, I clicked on some of these things. This one has verify compatibility with your memory configurator. I clicked on this, and entered a zillion different computers. Every consumer said this seemed to be a button that just says, can't verify compatibility. So no matter what you typed in, in some sense, it would just say always is not compatible.

So clearly, this page is designed to get us to buy this instead of buy that, and just telling you all the reasons. These words like OEM downgrade chips. I don't even know what that means. But I know it's bad.

It's not trying to teach us about the products. It's not trying to tell me why is CAS 2.5 latency better than CAS 3 latency. Or why is a six-layer board better than a four-layer board? Should you care about that? But it does seem like it's trying to convince us to buy this product and pay $25 extra instead of buying that product.

So what we're going to argue in this paper is that if you really had Bertrand competition for memory modules and they were being sold at cost, the retailers couldn't cover their fixed costs. The market wouldn't exist. The market does exist. Why is that?

Well, our thought is that firms are engaged. While Pricewatch has invested in technology to reduce search costs, the firms are fighting back with obfuscation to try to raise search costs and use those higher search costs to keep the increased markups. We can think of two different ways you could do that. One is-- the simplest theory of obfuscation is we know that you get higher markups and more dispersion when s is higher than when s is lower.

So maybe what firms are doing is just making s higher. If I make you spend 5 minutes clicking-- reading through-- unclicking the thumbscrews, unclicking other things, I spend 5 minutes clicking through doing all the things where I finally found out the price, I'm like, do I really want to go back and go through another 5 minutes with the next website, and find out what their real price is. So by just making shopping more time consuming, I've raised the s. Raising the s is going to raise the equilibrium price distribution.

And then the second thought would be-- this could be what I mentioned before, this add-on pricing situation where we know that equilibrium prices with add-on pricing are higher than equilibrium prices when all goods prices are advertised. Maybe what firms are trying to do is shift the equilibrium-- shift the-- change the game from a game in which firms advertise all prices to games in which firms advertise base goods, and then invent add-ons.

And that model, the larger is the add-on as a fraction of the price, the higher are the equilibrium markups going to be. So firms are in some sense inventing add-ons or inventing inferior products that the add-ons can be even larger. And that can be a way to raise equilibrium prices. And what we're going to try to do in this paper is look at demand and look at what demand looks like, and see, can this model add-on pricing explain the equilibrium markups that we end up with?

So the paper, we're going to look at four different products. The four different products are 128 megabyte PC100 memory modules, 100 megabyte PC133 modules, and 256 megabyte versions. We're going to treat these four categories differently, and just basically do it as an opportunity to have four separate products. We'll analyze each one separately, and just get some reinforcement of the conclusions we would get from one.

But we do spend most of the time-- I'll focus today on 128 megabyte PC100 memory modules. That was the product that was most popular and stayed alive for longest over the one-year data collection.

So what is our data collection? Within each of these categories, there's one retailer that we got internal data from. And then we have data, external data from all the other retailers from Pricewatch. The one retailer that we had this special data connection with, that retailer sold three prices-- three products within each category. We'll call them the low quality product, the medium quality product, the high quality product. You can roughly think of the low quality product as something that looks like this. The medium quality product looks like this. The high quality product looks like this.

Something to know about them is in this market, the medium quality products really didn't cost much more than the low quality products. Often if these cost $60 wholesale, these would cost $61 wholesale. The retailer we got the data from said, in fact, sometimes he would even sell someone the $60 product, and then ship the $61 product, just because dealing with the hassles and the customer-- customer disappointment when they got the low quality product, then it was broken. They returned it. It wasn't worth his hassles.

And you sometimes sell them a better product than they bought, just because these really didn't cost much more than these ones. But anyway, but he had the three products available. The high quality products that have the name brand chips in them did cost substantially more.

When you're thinking about this, though, note that high quality versus low quality, it's really a bundle. So it's not the case that you can compare my medium quality product from firm A and medium quality product from firm B, because quality is this many, many dimensional thing. What's the restocking fee policy? What's the shipping policy? How many layers does the board have? What's the CAS latency?

These are not something that's comparable across retailers. So we can really-- every different retailer may have its own low, medium, and high quality that don't correspond with the other ones. But so anyway, our retailer had these three different quality levels that were consistent.

The price data that we get is we went to the pricewatch.com website. And we just wrote an early web scraping program, and just basically downloaded the raw HTML files, and got the 12 or 24 lowest prices in each category. We did that hourly, and we did that hourly over the course of a year. Presumably, all the prices in that data are for low quality products. If the high quality products cost $15 more, they would never show up in the lowest 12 or 24 lowest prices.

We have quantity data from one retailer. So that one retailer we had contact with owned websites A and B. For those websites, we have the medium and high quality prices available. And we have the number of modules of each type that it sold. Typically, we're going to use this at a daily level. We're going to aggregate up to-- we have the real-time sales data, but we're going to aggregate up to the real-time sales, to the daily sales.

And then, also, something that's very rare is we have tremendous cost data here. So we actually have the raw-- we have all of the invoices that these-- from this retailer. So we know every time that it bought modules wholesale, what did it buy. And what's interesting about this business is this was basically a negative inventory business. This is a business where prices were falling over time. You didn't want to hold inventory.

And so what the retailer would do is basically every day would turn on its computers, and would start selling memory modules. And at the end of the day, or at 3:00 in the afternoon, it would say, I've already sold 12 of this module, and 24 of this one, and 36 of this one, and 14 of this one. And then it would buy them wholesale that afternoon, drive a truck up to city of industry in California, go to a wholesale distributor, and buy the 340 memory modules that it had sold that day online. Drive the truck back to the firm's retail office. And then the warehouse workers would then start putting those things in boxes, and shipping them out the next day.

So basically, as a result, it was buying things wholesale every day, or every two or three days. And we have very high frequency cost data, which is rare in understanding how the cost to the retailer varied.

Just for some basic picture of the industry, this shows the price of our 128-megabyte modules over the course of the year from which we collected data. So they start out at about $80. Prices rose to about $120. And then with technological progress, they drop. And by the end of the sample, those same modules that were $20 at the beginning or like $80 at the beginning are $25.

This is actually an interesting event here. This price spike was due to collusion by the memory manufacturers. So there was actually an antitrust case or legal case against these companies for price fixing. And companies were actually, I think, found guilty for price fixing. And they managed to raise prices dramatically, but only for a short time before the collusive agreement fell apart.

And actually, this is one of the cases, actually, one of the rare times I made some money on my research where I was contacted by a law firm, looking for high frequency data on these memory prices. And I just got this phone call out of the blue from someone who said, we're representing some company that's suing Micron or whomever about these things, and we wanted to get access to your data. How much will you charge for your memory data?

And I was like, charge for my memory data? I just give it away to everybody all the time. I haven't put it on my website at the time, but I normally just give it away for free. And they're like, how much do I charge? I don't know. $5,000? And the guy was like, yeah, sure. OK, thanks. We'll send you the check. And I was like, oh, I guess maybe I should have said $50,000. So I don't know. But anyway, these guys were price fixing. The price fixing scheme fell apart, and then memory prices continued on that downward trend.

What's notable about the downward trend for this paper is-- this is showing you, I think, the 12th-- the lowest price and the 12th lowest price for the memory over time. You can see that prices are changing dramatically. This graph here is our retailer's price. And our retailer changes its price occasionally, but as a result of changing its price occasionally, every time it changes its price, it may drop and be the lowest price in the market.

And then by leaving its price fixed, it's drifting, you can think, downward on that price watch list, where I leave my price fixed. And I go from having the lowest price to having the 10th lowest price. I lower it again to have the second lowest price, and then my price drifts across. And every time it drifts across as prices elsewhere are falling, I end up getting bumped up and down on that Pricewatch list.

So maybe I've just reset my price. And I set my price. And I'm here at $69. But then as I stay at $69, and this guy jumps to $68, and this one goes to $67, and this one goes to $66, I fall down the list, and then I reset my price. And I jump back up.

And we think that's very nice for identification if we're trying to estimate demand. We're basically regressing demand on your price. And we're going to think of the price, your relative price as exogenous variation. It's not that I'm deciding here to be the lowest price, and I'm deciding here to be the second lowest price. I'm trying here to be the eighth lowest price, and here to be the 10th lowest price. It's basically the firm's getting buffeted around by the randomness in prices. And we're going to think of that with where you are within the distribution as being largely random from the firm's perspective.

So some summary statistics, so for the 128-megabyte PC modules, again, we have two-- we have one year of data, but then we have two different websites. So days times websites, 365, we end up with 680 website day observations.

What do the memory modules cost? So the low quality modules generally cost, on average, $67. But obviously, remember, it was a big change over time. Medium quality modules are about $25 more. High quality modules are about $50 more. So very, very substantial add-on strategy here. The low quality modules are usually not very different from the lowest price in the market. The lowest price in the market is $63. This firm is charging $66.

What do people do when faced with these prices? Well, given that they are charging an awful lot of money for those add-ons, most people are choosing to buy the low quality product. They sell on average 12.8 per website per day. Medium quality modules are selling 2.4. High quality modules are selling 2.0. So most people buy the low quality modules. But a number of people buy medium and high quality modules.

Some people have also asked about, is this just some really tiny market? Well, what's going on is this firm has a website that has hundreds, and hundreds, and hundreds of products. So it's selling 17 PC100 128-megabyte modules per website per day. So it's selling 17 on website A. It's selling 17 on website B. But then it's selling many 256-megabyte modules, and many PC133 modules, and many PC whatever modules. So this is one of many, many products, but we just pick out this daily sales data, because this is the one that had things bought consistently.

The main thing we want to do in the paper is just analyze what is the demand curve. With all the obfuscation out there, competition between the sites, the attempt of this website to convince people to buy upgrades, what does demand look like and what do all the cross-price elasticities look like? We do this in a very simple thing. I think what we would like to do is, in some sense, we want to just regress log Q equals alpha 0 minus alpha 1 log P and so on.

You can't run a regression of log Q on log P when the Q's are discrete things that are often 0. So what we do instead is, in some sense, the-- you can think of it as log Q equals x beta plus epsilon. Instead of running log Q equals x beta plus epsilon, we just exponentiate both sides. We'll have Q equals e to the x beta times u, where u is a multiplicative shock with expectation of u given x equals 1.

What is on the right side? So we're regressing log Q on-- the log of the price of low quality memory, the log of the price of the medium quality memory, the log of the price of high quality memory, other covariates. Also, particularly what we've included to control for the sales is, what is the rank of your low quality memory on the screen? This turns out to be a very powerful predictor. If you're in the first spot on that screen, you sell much more memory than if you're in the second, or third, or fourth spot on the screen.

This is, in some sense, how we control for rivals' prices is I'm regressing sales at website A of medium quality memory on the prices at website A for low, medium, and high quality memory. And then I'm controlling for competitors' prices through these two variables. What's the lowest price on the screen? How high am I relative to that? And what's my rank on that screen?

So in some sense, you might want to include all the competitors' prices. That would, obviously, be infeasible. So I include, what's your rank on the screen, as well as my prices here. And notice that I'm putting low, medium, and high quality prices on the right side. Whether the left side has low quality memory as a dependent variable, medium quality memory, or high quality memory as a dependent variable. And I'm doing that because people are going to substitute between the firm's offerings. If it's charging $25, the upgrade is going to sell fewer upgrades than if it's charging $10 for the upgrade. It'll sell many fewer if it's charging $50 for the upgrade.

Our primary specification is just basically doing the OLS equivalent, treating the prices on the right side as exogenous for the reasons I explained earlier. Just expectation of u given the x is equal to 1. We do also in the paper have a couple different IV strategies. One IV strategy is using the costs as instrument for prices. Costs are a classic instrument for prices. So anyway, we use that wholesale cost data to get an IV.

The other thing we do is also use the prices for other products as instrument IV. The explanation for that one is if you imagine the firm comes in for the day. It sees that two of its warehouse employees call in sick, which is not at all an uncommon thing when you're paying minimum wage workers. If two of your workers are sick, what you do is you know, you can't package as many boxes. So you just try to sell fewer goods that day by raising your prices. You raise your prices, sell a few items at a high markup rather than a lot of items at a low markup.

And so I can use the PC133 prices as an instrument for PC100 prices, because if your employees are sick, you would probably raise both of them. And obviously, the PC133 prices aren't going to reflect demand for PC100, which is bought by separate consumers. But anyway, I think mostly we're going to say we think this is a case where within the distribution price variation really is as good as random. And we're mostly going to estimate this by, well, its GMM, but it's like OLS, where we're just treating those prices as exogenous.

What are some of the findings? So first, demand for low quality memory is very, very price sensitive. If your rank changes from 1 to 2 to 3, every time you drop down on that list, you sell many, many fewer units. We're going to report that price elasticities are as high as 33. So you raise your prices by 1% from \$100 to \$101, a third of your demand goes away. You raise price from \$101 to \$102, another third of your demand goes away again. \$102 to \$103, a third of your demand goes away. So it's just incredibly price sensitive demand for these firms caused by this rank effect. As I go from 1 to 2 to 3 to 4, I just keep losing customers.

The second thing is that-- interesting finding is there's a big loss leader effect here. So I'm selling this low quality memory basically at cost. But when I'm-- by selling this low quality memory, I get to the top of the list. And being at the top-- moving down the list-- when I move to the top of the list, I not only sell much more low quality memory. I also am able to talk some of those people into buying medium quality memory and some of those people into high quality memory.

So being high on the screen, I sell more of the product that makes my price high, but holding the price of the medium quality memory fixed, having a low price for low quality memory has this big loss leader effect. And I sell much more medium quality memory, and much more high quality memory when I'm up at the top of that list.

And this, again, it's controlling for the price of my medium quality memory itself. So it's holding that price of medium quality memory fixed. When I set a low price on low quality memory, I just sell much more of the medium quality memory because I'm drawing people to my site.

I mentioned that the reason add-on pricing can raise profits and equilibrium a few lectures ago is because of this adverse selection effect. When the cheapskates who come to your website to buy the low quality product won't buy the upgrades, you don't want those cheapskates buying stuff at a loss. And so that causes firms to have this incentive, flipping from undercutting to overcutting, and raises equilibrium price levels.

Here we're finding evidence of that adverse selection effect. So increasing this variable by one raises low quality sales by 129%, medium quality sales by 77%, high quality sales by 51%. So I'm selling more medium and high quality memory. My product mix is getting worse, and worse, and worse, and a larger and larger fraction of my sales are the low quality, low profit goods as I move up on the list. So this is the adverse selection effect we talked about, that it is clear evidence that adverse selection effect exists in this market.

In addition to these rank effects, the own price effects are significant as well. So when I raise my price for medium quality memory, I do sell less medium quality memory. When I raise my price for high quality memory, I do sell less. Those elasticities are like 5 or 6. So if I raise my medium quality price by $1, I sell 5% less. I raise it by $5. I sell-- sorry-- by 5%, I sell 25% less. So this does appear to be substitution within the website, where if those upgrades are 15, 30, or those upgrades are 30, 50, I sell much fewer when I have higher prices.

So when we add together these two effects, the rank effects and the own price effects, this is giving us the effects of the different categories. But I'll focus on the 128 megabyte PC100. What I get is this incredibly elastic demand for the low quality memory, and then somewhat elastic cost for the other memory. I get these two notable off-diagonal elements saying that there's that loss leader effect. The other off-diagonal ones down here are positive, as you would expect. When I raise my price on medium quality memory, I sell a little bit more low quality memory, because we'll substitute away from the medium quality memory. I'll sell a little bit more high quality memory. But again, these are small effects.

I said that. Those are the loss leader effects. I also said that. Anyway, so we've got this-- again, just summarizing this super elastic demand for the low quality memory and price sensitive, but not nearly as elastic demand for the medium and high quality memory. Naively, this is going to make you think that these products are going to be sold at higher markups because demand is less price sensitive.

Question now-- I think the paper's biggest contribution is saying, we know that this add-on pricing theory is correct. If you sell products, if you have-- get people to compete by only advertising low quality products, and then having to-- selling the add-ons at high monopoly prices, we're going to get higher equilibrium prices than we would get without that.

The question is, how large is that effect quantitatively? Does that affect-- is that effect large enough quantitatively to explain what's going on? And in fact, how big are the markups here? Are they explained by this add-on pricing theory?

To do that, the paper has a generalization of what was in the QJE paper, noting that a formula for price cost margins-- if you remember, standard single good price-- standard single good monopoly pricing is price minus cost over price equals minus 1 over the elasticity of demand.

What I note here is that there's a generalization of that model, of that formula to a situation where you have a firm selling both a base good and an add-on. And what this is, so this is a generalization of the price cost margin. So this is the price for low quality memory plus the x is going to be the fraction who upgrade when you're competing against another firm, and your price is p1L. The other firm's price is p2L.

So this is price plus fractional upgrade times the margin times the upgrade price. You can think of this as what's your equilibrium price per customer, taking into account that some pay more than they think they're going to pay when they click on that link, minus the cost of providing that good, which is the cost of low quality memory, plus the fractional upgrade times the cost of buying the upgrade.

And then, again, this is the price. So it's just price minus cost over price, where in price, I've included the fraction upgrading times the upgrade price. So this is per consumer amount paid. What you get is that price minus cost over price is minus 1 over epsilon times 1 plus something. And I sometimes refer to this as the adverse selection multiplier.

The worse is the adverse selection multiplier, the higher is your markup relative to 1 over epsilon. The key term here is this one, which is, what's the upgrade profit times how does the fraction who choose to upgrade increase when your price increases? So as you get further down on the screen, have less of an adverse selection, how much more successful are you at selling upgrades? And this formula predicts that as the adverse selection problem gets worse, the dx / dp is bigger. And the p upgrade minus cu is bigger. You're going to have a larger margin relative to the minus 1 over epsilon, which is the naive level you would expect from the single good intuition.

Anyway, so the paper shows that this is true if you define epsilon to be this thing we call the overall elasticity, which looks an awful lot like a derivative of low quality demand with respect to low quality price. What we do is-- again, this paper is unusual in that we have this cost data. So we have the daily cost for each of these modules.

We also know just from talking to the proprietor of the firm, what does it cost-- it's not rocket science. What does it cost to hire a minimum wage worker to take the memory module off the shelf, put it in a cardboard box, put one piece of tape on it, run it through the postage meter, and send it to the customer? And so the firm was able to give us very good estimates for what is the cost of sending a memory module-- of packing a memory module into a cardboard box and sending it.

So we're able to estimate markups here. And what we find is that, qualitatively, these loss leaders, the low quality memory, it's being-- it actually is being sold at a loss most of the time. So here it's approximately zero markups. But low quality memory's average markup is minus 0.7%. Medium quality memory has got a 17% markup. High quality memory has got a 27% markup. So overall, when you add up the 12 units of this, two of this, two of this, the average mark up is 7%.

How does that compare with what you would have expected had this add-on strategy not been going on? Well, the overall elasticity is minus 23.9. So if you had the simple price minus cost over price is 1 over epsilon pricing, you'd expect a 4.2% markup. And what we see is 7.7. Can we explain why we're seeing 7.7 instead of 4.2?

Well, if each of the things in this adverse selection multiplier are also things that we're estimating our demand system. So when I estimate those, I find that my estimate for the adverse selection multiplier is two. So when we estimated this demand system without using any supply side information, without imposing that the firms were maximizing profits, and yet what we find is that when we look at the estimated parameters, with the estimated parameters, the equilibrium markup prediction is almost exactly the true markup prediction.

So this is saying that, we estimate the adverse selection multiplier is two. And that does seem to explain why profits margins are 7% instead of 4.2%. We do this across several different categories. It predicts that the adverse selection is even more important for PC133 memory. It predicts that it's about the same for the 256 megabyte memory. And in each case, we find that the 1 over epsilon formula would say you'd get 3.6% markup. We actually get an 11.5% markup. But that's perfectly explained by this adverse selection multiplier that says, yes, it should multiply markups by 3.5%.

So I think the basic takeaway from this paper is that one technique that firms are using for obfuscation here is, well, one technique is making things annoying. That's going into what determines these-- what determines the sensitivity to rank and how quickly people give up. The second thing that's going into margins is creating this adverse selection effect by offering add-ons. And the amount by which markups go up relative to the single good model seems to be almost exactly in line with the adverse selection effect.

So it seems like, in some sense, we don't need anything other than the adverse selection effect and the level of search costs to explain what's going on here. It does seem like the adverse selection has an effect. It has an effect as big as we would expect it to be. And it's not a puzzle why markups in this market end up what they are, which is markups are quite low, but not so low that these firms, which are bare bones firms with no advertising, can't survive.

And I think the rest of these slides just have the same things I just said. So yeah, here's where you see the low quality memory has got negative, and the others have positive. There's the adverse selection effect. There are the predictions.

And then this summary of the conclusions that I told you that observed markups are very close to ones you'd predict from the demand estimates. Adverse selection is roughly doubling average markups. And markups for low quality products are very low.

So in my final 15 minutes, I'm going to give my final paper, which is Galenianos and Gavazza's paper about "Regulatory Intervention in Consumer Search Markets, The Case of Credit Cards." One thing that I often do in this course is tell students that it's a mistake to be too ambitious when you're trying to come up with paper topics, or you don't need to be so original and creative. It's a great idea to just take someone else's paper and improve it, because if you take someone who's at the state of the art, you improve what they're doing, you know you have a state of the art paper, whereas if you start from scratch, you're going to end up-- you're going to end up very much behind the frontier.

Normally, this is advice for writing theoretical papers, because you take a theoretical paper. You can just-- it's a well-known thing. You take a theoretical model. You improve on that. You've got a new and better model.

The thing I want to point out here is one can do that in empirical work as well. One of my earliest published papers was a paper where I took-- you'll see later in the semester where I took a published empirical paper by Rob Porter. I just did some additional things on it. That was one of my first published papers.

Galenianos and Gavazza decades later are actually doing the same thing here, with the idea that they must have read Stango and Zinman's paper, and they thought Stango and Zinman is a very nice paper, showing that search costs seem to be related to markups. And what they're trying to do is, say, Stango and Zinman said that search costs are related to markups. Can we actually turn their paper into a structural, quantitative paper, where we say, can search costs explain the magnitudes of everything that they're finding in their data? Or how would we say, is their data really consistent with the search cost explanation versus another explanation?

It's also, I think, a very nice paper in that it's doing something we don't do often. It's doing something different in IO. Like many people have this view that every IO paper has to have this really, really complicated individual level data set and estimate this sophisticated model observation by observation. And what Galenianos and Gavazza do in this paper is they're basically doing a macro style calibration.

And I think it's not what other people would have done with this data, but I think that it really makes the case that doing a calibration like this can be a very interesting way to think about whether a theory really explains the data or not. And do we always need to do the same kind of structural estimations we do? And I think it's a very elegant paper. It also makes some nice technical observations. And so anyway, I'll do it quickly here, because I don't have much time.

So let me say, the basic thing is what they show in the paper-- think back to what Stango and Zinman showed. What Stango and Zinman showed is that consumers get very different offers in their offer database, and that consumers end up paying very different amounts. And so we have these sort of two data sets. We have this offered data set, and then we have the sort of what consumers-- the accepted offers, what consumers end up paying.

And the starting point for Galenianos and Gavazza is think about we have these two separate data sets. The offers are going to lead to accepted offers based on consumer preferences. So if you think about what consumer preferences look like, what are their search costs, how much they view these offers as differentiated, that's going to lead us to some mapping that maps offers that they receive into what offers they accept. And notice that you can think about two different ways to explain the mapping.

We know that these offers are quite dispersed. And we know that these accepted offers are fairly dispersed and quite high. You could explain this mapping in two different ways. One way would be that consumers have search costs. Therefore, they don't open a lot of the envelopes. Therefore, they don't all take the lowest offer. And therefore, they pay high prices, or they pay somewhat dispersed prices, because some of them open the low envelope first and pay that one.

The other way to explain this same fact would be with consumers just have very idiosyncratic preferences. They get these different offers. They open all of them. But they care a lot whether their credit card says Bank of America or Capital One on it. Therefore, some of them pay much, much more than they need to, even after having opened all the envelopes, because they really like that credit card company.

[INAUDIBLE] note is that we really can't just from thinking about this mapping tell the offer, the differentiation explanation, apart from the search cost explanation. But there's also a mapping that goes in this direction, where if you know what consumers are going to accept, the firms who are profit maximizing should be choosing their offers optimally. And if consumers had these very strong idiosyncratic preferences and viewed the offers as very differentiated, then all the firms would be sending high offers. Whereas if it's the search cost explanation, the search cost explanation could perhaps better explain why consumers are sending low offers.

That's sort of intuitive, that we need the search cost to get the firms be setting these low and heterogeneous offers. But Galenianos and Gavazza are trying to say, quantitatively, is this true? Can we say that a model with search costs can explain the data quantitatively, whereas a model with differentiated products cannot? So what they do is they write down a model that's simple enough that they can estimate what the equilibrium is for a set of parameters. And then they say, we're going to calibrate the parameters to match some macro moments in the data. And then see if we can fit the data.

So what does the model look like? So anyway, the model has buyers who have value z. z is distributed according to some distribution m on some interval. So this is your value of having a credit card from any firm whatsoever. In their model, some customers will have z relatively low-- won't get a credit card, or won't buy any offer.

There's some lenders who have mass l. These lenders have heterogeneous costs of supplying funds, k, distributed g on k lower bar, k upper bar. As I said, these heterogeneous costs are going to be the way they generate the price dispersion in the data. In the Stahl model, you can generate price dispersion by-- generate price dispersion by all the firms are exactly identical, and they just randomize over prices. But there's an earlier literature on price dispersion, going back to Jennifer Reinganum has one of the earliest models of this variety, where you can also get price dispersion as a pure strategy phenomenon by simply allowing the firms to have asymmetric different costs.

And then where the firms with identical costs randomize, you can have a purification of that, where the firms with low costs set low prices. Firms with high costs set high prices. So they're going to add this lender heterogeneity. And then, interesting model here is that rather than having buyers choose an equilibrium, how many offers to get, the buyers are going to exert an effort level s. And then if you exert effort level s-- you can think about I'm going to spend one hour on my computer, looking up credit cards online, and see what offers I can find.

They have a Poisson random number of offers. So if I spend time s searching, I get s of-- in expectation, I get s times L offers, but I get a Poisson random number of offers. So I may actually find zero offers. I may find one. I may find two. I may find three.

Doing this effort has a nice feature that s is a continuous variable. I can do a first order condition for what's the optimal amount of search effort s to make. You might think, well, yes, you made s continuous, but now you've got this nightmare, where instead of people having-- all having one offer, or two, or three, you've got-- I'm getting this Poisson random number of offers. How am I going to do that at all?

An observation in this paper is that-- and there also been some others who have recently have done the same thing, is that this Poisson random number of quotes can actually be a more tractable model than having a deterministic number of quotes. And so because it's got this feature that effort is continuous, giving us a continuous first order condition, and the expected value to search is actually tractable, this can be a very nice framework. And I think other people are going to want to do this in the future.

And let me say, also, if you remember back to intuition for Stahl, to get heterogeneity in accepted offers, what Stahl did is have some people who get very-- some people who have negative or low search costs, who search every firm. Some people with high search costs search all the firms.

In Galenianos and Gavazza, they don't have randomness in the search costs, but what they have is this randomness in the value z. If you have a high value for a credit card, you're going to want to get a lot of quotes, because you know you need to buy at least one. If you have a low value z, you're going to do a little search effort, because you're going to be perfectly happy to not buy from anyone. And it's the heterogeneity in z that's going to lead to heterogeneity in search effort. That's going to lead to equilibrium price dispersion.

Anyway, so going over the model quickly, buyers have this payoff from getting a card, which is their raw value from getting a card minus the disutility of the interest rate minus an epsilon, which allows for this sort of we're going to have the horse race between search costs versus product differentiation. So they allow for product differentiation.

They write Fc for the CDF of a random draw of this plus this. That's obviously something that's determined in equilibrium. And then the seller's payoff is going to be the probability of-- the probability that someone buys from you if your interest rate is R times the payoff you get if your interest rate is R, which is R times 1 minus rho, minus k, which is the cost of capital.

And I see I've got a misprint here. R-- rho would be the default probability here, because your payoff is R times the probability the customer pays you back, which is 1 minus rho. So rho would be the default probability.

Anyway, as I said, the Poisson search model is surprisingly tractable. Effort is a continuous variable. And then when you think about what's the probability that a consumer accepts an offer with total cost c, the probability that a consumer accepts an offer with total cost to them of c, you're summing over the number of offers they receive. n equals 0 to infinity. So n is the number of offers they receive.

If they receive n offers other than yours, they're going to buy your product with probability 1 minus F of c to the n-th power. They only buy your product if every other product they get has higher cost to them. Or again, cost includes the financial cost and the idiosyncratic preference cost. So it's the sum of e to the minus-- so this is the Poisson probability that if they do effort, I guess if they do effort, a sub z, alpha sub z, they get n costs that's e to the minus alpha z, alpha z to the n over n factorial.

If you look at this, it's this sum. And you notice it's got e to the minus alpha z over n factorial, alpha z to the n, 1 minus F of c to the n. So it's just like this and this go together. And so the probability they buy from you is just e to the minus alpha z times F of c. And then when I set a price, I'm setting an interest rate of-- I'm not setting a cost c to them. I'm setting an interest rate, and then they have an epsilon. So there is an integral over the epsilon to determine the probability of purchase.

But anyway, we get this sort of nice closed form for the probability of purchase as a function of the interest rate that I do charge on a credit card. And so this sort of like-- in some sense, this tractability of the model is just a key step that makes it possible to solve the model given any set of parameters.

What do they do? Well, they make parametric forms for many of the primitives. For example, assuming that the buyer's value distribution is log normal with parameters to be estimated. I guess this is off a working paper version that had 11 model parameters and 15 moments. I think the published version of the paper has-- referees always ask you to do more. I think the published paper has 24 model-- 28 model parameters, and 36 moments to be matched.

But basically, you're setting model parameters-- parameters of the various distributions that are there to match a bunch of moments in the data, like, what is the 10th percentile of the accepted offers, median accepted offer, 90% offer? What fraction of consumers are getting at least two offers? What's the median number of offers conditional on getting two?

So you're trying to explain simultaneously what are the accepted offers here and what are the offers being sent. We're not trying to explain them customer by customer. We're just trying to say, I want a model that gives this amount of distribution that's observed in the offers and the amount of dispersion that's observed in the accepted offers. And I need to match other parameters that I know from the data.

Anyway, so what they do is they estimate these things. And then the-- I won't go through it, but the main thing they then have is with this calibrated model, they talk about a counterfactual. Counterfactual is, what would be the effect of capping credit card interest rates? So many people pay 35% interest on their credit cards. You could ask, suppose in instead of charging everybody 35% on the credit cards, you put a law that says you can't charge more than 27.5%, or 25%, or 22%?

That's going to have two effects. One is that some people are going to pay much less on their credit cards, because they're going to get an offer that has a lower rate in the mail. They're going to buy an offer that is a lower rate. Some people are going to be hurt by that, because they can't get a credit card, because people would not be willing to lend to them at the lower 25% rate. They're just going to get no credit card at all.

What they find-- do I have things here-- is that they find that actually here, their model predicts that laws capping credit card interest rates should work well. What we get is that when you cap interest rates at 27.5% with these parameters, that's good for consumers. Yes, it is true that some consumers are locked out of the market, but the number of consumers who are locked out of the market is fairly small, and many other consumers benefit from the drop in interest rates.

They do notice that-- they do have results showing that these are not monotone effects. Lowering the-- what's best for consumers is not to have a cap, obviously, of zero, because then no one would be served. And in fact, in their data, I think they show that consumers are better off with a 27.5% cap than with no cap. They're better off still with a 25% cap. When you go to 22.5%, consumers start to be worse off.

So anyway, I think it's a very nice paper in two dimensions. It's got this sort of-- it's got originality, and thinking of doing a calibration, rather than doing what we normally do in IO. It really makes the case that this model quantitatively can explain what goes on. And several of the key moments can be explained simultaneously by this model. And then it has this sort of counterfactual, saying that, here's a policy relevant question we care about. What would happen if you cap interest rates? And with their model, they predict that capping interest rates at still high, but moderate rates would be good for consumers.

So that's all I've got for you. So I guess what we're going to do, we have a bit of a non-standard timing in the physical class at MIT. Jean Tirole is going to come give guest lectures next week. So we'll have two lectures from Jean, and then Tobias will come back the week after and talk about structural work on search. I guess in the OCW course, for those of you who are watching this there, I think we may flip back to the natural order, and have-- next lecture will be the third lecture on search, and then we'll have the guest lectures by Jean on platform markets.