[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK. Let me go ahead and get started. So today, I'm moving on to the next unit on the syllabus on consumer search. And if you think back to what I was doing for the previous week and then what Tobias did last week, it's following what is very common in IO today, which is to put a tremendous emphasis on product differentiation as the source of markups in markets.

And I think it's just helpful to think about, when we do that, we're both explaining markups using product differentiation and also, whenever we're making welfare calculations, we're attributing people's-- the product differentiation to people's underlying true preferences for goods. And we're assigning-- a significant part of welfare is then, what are people's epsilon ij's, and what is the average epsilon ij? And that's considered part of the utility function. That's something we consider when we're doing welfare.

And I guess it's interesting. It reflects current IO, and it also reflects historical IO. If you look at Tirole's IO textbook, Tirole does not have a chapter on search. Tirole exclusively does product differentiation and repeated games as sources of markups.

But intuitively, when you think about products-- and a lot of important products people buy today, it's unclear whether there really is much differentiation or whether the differentiation is at all sufficient to cause what-- the outcomes we are seeing in the market.

So one example is credit cards. How many people have a really strong preference for whether they have a Capital One credit card or a Citibank credit card or a Bank of America credit card? Is it that I really like the color of the Citibank card much more than I like the color of the Bank of America card? Or is it the automated robot customer service agent I get with Citibank? Is that much better than the automated robot customer service agent-- I like the voice of this one, A versus B? It's very hard to think that that's what's driving the markups.

Cell phone plans-- Verizon and T-Mobile will spend a lot of money having commercials with people talking about how great their service and reception is. But again, how many people have very strong preference for, I really like the Verizon service on my phone much more than I like the T-Mobile service on my phone? Mortgages, again.

And maybe one other story. If you think about even goods where there's clear natural product differentiation, is that product differentiation really explaining the markup? So, for instance, I own a Hyundai Kona, which is a little, compact SUV. Let's suppose you-- you guys live in Cambridge or whatever, or in Somerville on top of Cambridge. Suppose you want to buy a Hyundai. What can you do?

It turns out, there are no Hyundai dealers in the city of Cambridge. Instead you can go up to Arlington, you can get one. You could go out to Lynn, and you could get one. You could go down to Quincy and get one. And each of those places is, like, half an hour away from you. I guess Arlington is probably the easiest from Cambridge on public transport.

And if you have a car way out here, there's dealers in Norwood and dealers in Framingham that you can get to on the highway. So for me, I could get to any of those car dealers in half an hour.

And so if you think about it, the Hyundai Kona bought from the Arlington Hyundai dealer versus the Hyundai Kona bought from Boch Hyundai in Norwood, what's my preference for buying from dealer A versus dealer B? Again, it is the exact same car offered at both dealerships.

If you think of it like, OK, you can think of this as geographic differentiation, where, for me, I'm equally good driving here or driving here. But there are some people who live right in Arlington for whom that is an annoying drive. But if you think about the p star equals c plus t model, what is t? t is the distance by which someone at one end of the line prefers their product to the people at the other end of the line.

You can do that drive in 35, 40 minutes. What's your disutility of driving 35 or 40 minutes one time to go get your car or something like that? Is that $20? Is that $30? If you're a high-income person, is that $50 or whatever?

But that would then say the prediction is that all these cars, the Hyundais at all the dealers, should be marked up by $40 or $50 over the dealer cost to the-- and we think that that's just-- counterfactually, that's just not correct. Buying a car, local dealer markups are $1,000 or something like that. And the c plus t model is not going to give us that.

So how do we explain that? One idea that people have is that another important component of real-world markups could be search costs. And search costs, intuitively, the more complicated or difficult it is to find out what the prices are, the more people are going to give up without finding what the prices are, and the more that prevents the competition that Hotelling says should be there for those marginal consumers.

In the Hotelling model, if you estimate that people-- that there's $1,000 markup-- let's suppose that you're doing a differentiated product estimation and you're imposing the supply side. I think Tobias pointed out that sometimes you just do it entirely with demand-side moments. Sometimes, you also impose the supply-side moments to explain what's going on and use the cost data.

Well, if you're imposing the supply-side moments, you're saying there's a $1,000 differentiation between this Hyundai dealer and this Hyundai dealer, then you would get-- it would be terrible if one of those Hyundai dealers closed because all those people are getting $1,000 worth of utility from going to this one versus that one would have that huge loss in welfare, which we think is probably not a correct answer.

So anyway, I'm going to start this with-- Diamond has this very nice paper from 1971, noting that this basic idea that the higher is search costs, the higher price is going to be seems harder to capture than you might have guessed.

So here's his model. You have N firms. They're producing a homogeneous good, like the Hyundai Kona. Oh, no. Actually, it's not a Hyundai Kona. So this is going to be credit cards. So N firms produce a homogeneous good at cost c. So that cost c is it costs them 8.99 or 9% to loan you money on a credit card. And so c is 9%.

They're a continuum of consumers, and they have identical multiunit demands D of p. So if they give me a credit card at 9%, I may borrow $10,000 on it. If they give me a credit card at 18%, I may borrow $3,000 on it. But there's some downward-sloping demand function for how many units you're going to buy from the firm as a function of the price. And I'm doing the credit cards rather than cars because it's important that you're buying more than one, whereas cars, typically, you buy exactly one.

I'm going to assume that p minus c times D of p is a concave function. It's got some finite monopoly price pm. And I'm going to consider the following game.

First, the firms simultaneously choose their prices. And then customers have to pay s every time they get a price quote. And we're going to want to assume that they search optimally, get some number of price quotes. And then after getting some number of price quotes, they buy D of p units from the lowest-priced firm that they found. And we're going to want to have some optimality in this-- search is somehow optimal, given how much the cost of every search is going to be.

One immediate subtlety is we need some new solution concept other than Nash equilibrium. Because in a Nash equilibrium, everyone behaves as if they know the strategies of the other players.

And so if you think about this as an N plus 1 player game, there's N firms and one consumer buying, then in Nash equilibrium, you're implicitly assuming that the consumer knows the firm's strategies. And if the consumer knows the firm's strategies, their strategies are their prices. And the consumer can just buy from the lowest-priced firm and not have to pay the s to find that firm.

So what do we do? Diamond, essentially, in his paper, made this an N plus 2 player game. So you have the N firms, you have the one consumer, and then you have the one player I'll call "nature." And what the firms do is-- so first, stage one is the firm's-- they choose p1 up to pN. So that is the first thing-- actually, I need a timeline.

So p equals 1, the firm's move, and they choose p1 up through pN. And then at stage two, nature moves. And what nature does is chooses a random permutation. So it's if nature builds a room, and this room has doors numbered 1 through N.

And so the firms pick their prices, and then nature takes firm 1 and sticks it behind a closed door. Say it randomly picks door 7. And then it takes firm 2 and puts it behind another door, say, door 1. And so nature does this random permutation. And then t equals 3, the consumer starts opening doors. And each door you open, there's a cost of s.

So in equilibrium, I'm going to know what p1, p2, and pN are. But I don't know which firm is behind which door. So then I open door number 1, see who's there. If I don't like that price, it seems high, I may open door 2. If it seems like it's still worth paying s to open another door, I'd open another door, and I continue like that.

And so then Diamond asked, what's the equilibrium of that model? And the thought was that you might get an answer that says, when search costs are higher, prices are higher than they would-- without the random permutation, this is a Bertrand competition model. And so you might think that you would get a just Bertrand outcome, everyone's p star equal c. But then as search costs go up, maybe prices go up in equilibrium. And they do, but they do more than Diamond wanted.

In some sense, he started this literature with this thing that sometimes people call the "Diamond paradox." If you take any search cost s that's strictly positive and less than the consumer surplus, the monopoly price, then any perfect Bayesian equilibrium in which any consumers search has all firms pricing at the monopoly price.

So in some sense, the only equilibria of this model are it's a discontinuous jump in prices. As soon as there's any search cost, even a penny, the only equilibrium is everyone charges a monopoly price instead of everyone charges c. So it says, Bertrand competition is an extraordinarily fragile result. Bertrand competition relied on zero search costs. You put any search cost in, you don't just get a little markup, you get the monopoly markup.

Now, we don't believe this, but I think it's good to see why this happens. So first, the argument that this is an equilibrium is straightforward. Let's suppose you have a world where all the firms charge pm, and you know that everyone charges pm. You've grown up. You've lived your life in this world. You know that, yes, everyone charges a monopoly price.

You open the first-- you get assigned-- you have to play this game. You open a door, and you're like, oh, what does cell phone service cost? Oh, it's $59.99 a month, even though it has marginal cost of zero for them to provide it. What do you do? Do you keep opening the other doors? You're like, no, every cell phone company charges $59.99 a month. Why keep opening doors? I open the first door, I see the monopoly price, I pay it, and then I'm done.

And could a firm cut its price and gain consumers? And the answer is no, because when consumers expect everyone to be charging the monopoly price, everyone stops at the first door and just buys from that firm. And so suppose you have a lower price-- well, then consumers-- and you get assigned to room 7.

If a consumer comes to you first and opens your door, they'll be like, oh, bonus, I get to pay $35 a month instead of $59.99 a month. But it's only the consumers who open that door first who buy from you. And they would have bought from if you'd charged the monopoly price. So the firms don't get any extra consumers by charging a lower price.

If everyone charged a monopoly price, consumers don't search. If no one searches, all firms charge monopoly price. That's an equilibrium. The less immediately obvious part of this argument is that there's no other pure strategy equilibrium.

To think about that, just imagine we had some other equilibrium and the firms charged different prices. So maybe this is the monopoly price. And then some firms charge a monopoly price, but some firms charge less than the monopoly price.

Well, the argument is going to be, you can't have an equilibrium like this where somebody charges a lower price than everybody else. Because if you're charging a lower price than everybody else, I could raise my price by epsilon. Again, anyone who opens the door and sees my-- so consumers were expecting that there was some firm that charged $35 a month. And I go to $35.50 a month.

So consumers open that door. They may think, oh, no one was supposed to charge $35.50. There was supposed to be a firm at $35, a firm at $36, and a bunch of firms at $40. But I just found $35.50. What do I do?

Well, the most you could expect to save by opening another door is $0.50. Because you thought that there was someone at $35. You didn't find the guy at $35. But that's the best you're going to think to do.

As long as I keep the amount, I raise the price less than the s, anyone who opens my door first is still going to buy from me. So I don't lose any-- anyone who ever opens my door is going to buy from me. So therefore, I can't lose any consumers by this price range. Therefore, I'm always going to raise my price by epsilon over what people expect.

And so that makes it impossible to have an equilibrium where somebody is below everybody else. It also even makes it impossible to have an argument where you have two people tied for the lowest price. Because again, I just make an epsilon increase. Anyone who opens my door, as long as I'm within epsilon of the lowest price in the market, they're all going to buy from me.

So it's going to have to be that you can't have anyone-- the lowest-priced firms are always going to want to raise their price by epsilon if they're less than the monopoly price. Similarly, I can't have anyone charging greater than the monopoly price. Because if they're charging greater than the monopoly price and not selling anything, they make no money. If they are selling things, they'd be better off going down to the monopoly price itself. So the only equilibrium of this model is everyone charges monopoly prices.

So I think it's a useful thought experiment, but it's obviously not what we want to do in practice. And so then where the search literature has gone is thinking about, what happens-- what can we do with search to get search to have a continuous effect like we think it probably does-- should have? And what's wrong with this argument?

One way in which this literature has gone, and which is completely realistic, is models with heterogeneity in search costs, where there are some consumers who have higher search costs and some consumers who have lower search costs. And it turns out that-- it's a nice literature because it actually turns out it's going to explain two phenomena simultaneously.

One is it's going to give us models that do have positive markups that are continuous in the search cost. And then the other thing it gets us is price dispersion, where different firms charge different prices. And I think in the car dealer example, you do expect that every different car dealer has a different price that they're willing to settle for on the car if you negotiate long enough and that every different firm is going to charge you a different price.

So how would this model go? So again, let's stick with my credit card example. So I have N firms. They're all producing that homogeneous good at the same common, constant marginal cost. Again, I'm going to let them choose prices. Again, I'm going to have the same continuum of consumers. Each consumer has a multiunit demand D of p.

And I'm going to have consumers searching optimally, as in the Diamond model. But instead of the Diamond model where everyone had search cost s, I'm going to have multiple types of consumers. I'm going to have a fraction mu of the consumers. Sometimes people refer to these guys as "shoppers."

Shoppers have a negative search cost. So shoppers actually get utility out of comparing prices at different firms. And you may think that's unrealistic. But I mean, there are some people who like shopping as a hobby.

There are some graduate students who love to tell me that, I booked this crazy plane ticket where I fly to here, and then I have a return ticket that's booked out of this city. But then I'm going to get off in this city and save that thing. And I'm saving $7 by doing this.

And you have to think, did you spend, like, 14.5 hours figuring out some really clever way you could save $7? It must be that you got a great deal of utility out of figuring out all of the different ways one could fly from A to B and feeling that you were-- whatever. I don't know.

But there are some people who just will enjoy shopping and will do more shopping than seems reasonable. And so I'll model them by having this negative search cost. And then I'm going to have a fraction, 1 minus mu of consumers, who have the more traditional positive search costs.

And at least at first, I'm just going to assume that these people have search costs in some interval s lower bar, s upper bar. It could be distributed with some density. But they are all less than the consumer surplus, the monopoly price.

The reason for putting this upper bound on search costs is, if your search costs are less than the consumer surplus of the monopoly price, then you run into this problem of just those people aren't going to search at all. And then they're not going to search at all, so then no market exists.

So I want search costs low enough so that consumers are going to want to search and buy, even if the prices are high. And I don't have to worry about that edge case of what happens if prices are high-- too high.

So this distribution of search costs. So I've just got some interval s lower bar, s upper bar, where I have a density of consumers. And then you can think of it as a point mass of consumers with negative search costs. And it's just a simplification, but it's a simplification that implies that these consumers, it's easy to solve for what they do. What they do is they visit every single firm and then buy from the lowest-priced firm that's out there.

So what's Stahl's result on this model? Part a of the theorem, this model has no pure strategy equilibrium, but it does have at least one symmetric mixed strategy equilibrium. And in this symmetric mixed strategy equilibrium, the firms choose prices from an atomless distribution F that has some support p lower bar, p upper bar. And that support does not contain marginal cost.

So what's going to typically happen here is, if you look at pricing-- did I give it a name? I did not give it a name. So we're going to have this distribution where there's-- you could price it marginal cost in Bertrand, but firms don't. Instead, there's some level p lower bar, and there's some level p upper bar.

And then what firms do is there's some density. It might look like this. And so what firms do is they choose prices-- use a mixed strategy, choose prices from a density that looks like this. And so prices, from the consumer perspective, are random.

And then, in some sense, this mixed strategy on the firm's part then rationalizes the search on the consumer part. Because if the firms are mixing over prices, you go and find a firm's price. What you're getting is one draw from the mixed strategy distribution. You may get a high draw. You may get a low draw. You may get a high draw and decide to search again. You may get a low draw and decide to stop your search.

Clearly, if you draw-- if the first firm you draw actually has price p lower bar, you're always going to stop there because you know that's the best you can do. If the firm has price p upper bar, if your search cost is really high, like if your search cost is near s upper bar, you're going to buy-- the model is going to be such that you are going to buy from that price if you find it. But then someone with a lower s may decide not to buy there and only buy if they find this price or better.

So, Remarks. Yeah, the mixed equilibrium can be seen as a potential explanation. Yes, question.

AUDIENCE: Yes. So intuitively, why do you need the shoppers in this model? Because at first glance, it would seem like just by having the dispersion in search costs would generate a dispersion in prices.

GLENN ELLISON: Yeah. Unfortunately, you get this same-- it's the same Diamond argument would work here. So let's suppose that we had these-- I took this away. And then imagine I tried to have an equilibrium like this. And then think about the firm that's charging price p lower bar.

So to have a mixed strategy equilibrium, I have to be indifferent over all these prices. Let's suppose I deviate, and instead of charging p lower bar, I charge p lower bar plus epsilon. Again, the same argument.

Any time anyone gets a quote from me and they find price p lower bar plus epsilon, as long as consumer surplus at p lower bar plus epsilon, consumer surplus at p lower bar minus consumer surplus at that is less than s, they're going to have to buy from me. Because the most they could gain by finding p lower bar is this. If this is less than s, they would stop.

So you still get this epsilon over a cutting argument that says, this can't be an equilibrium, because if everyone knows that everyone charges at least p lower bar, then everyone would want to charge at least p lower bar plus epsilon. And if everyone knows everyone wants to charge at least p lower bar plus epsilon, then everyone wants to charge at least p lower bar plus 2 epsilon and the thing collapses.

AUDIENCE: So you're going to want someone who's always willing to keep searching.

GLENN ELLISON: So you want someone who is always willing to keep searching, that you may lose-- you need to explain, how could anyone ever charge the lowest possible price? And there, it could be, if you raise your price by epsilon, there are those shoppers who you're going to lose by doing that.

AUDIENCE: So in this model, the consumers always know the distribution?

GLENN ELLISON: So in this model, we're assuming the consumers know the distribution. That is, this is where-- obviously, you can have models where consumers have to learn the distribution. But here, we're sticking with the equilibrium notion of consumers are best responding to the firm's strategies, and the firms are best responding to the consumer strategies. Yeah.

And I think that one doesn't strike me as so unreasonable that, like, on the credit card market, I know there are credit cards that offer a good deal and credit cards that offer a bad deal. And I go out there and I start clicking on credit cards, and I look at this credit card and I see what rate they're offering.

And certainly, the credit card offers you get in the mail, they all have different rates on them. I kind of know what the distribution is, and I know what's going to take me-- I get some sense of how good I'm doing by following it.

And like we said, they know the equilibrium strategies. They don't know the actual distribution. So if a firm deviates, it changes the distribution and consumers don't know that. So consumers know what's the normal distribution of prices. Yeah.

So it's a mixed equilibrium. You can think of this as explaining why prices differ across firms. Prices in this model move continuously with search costs, and they move continuously with the fraction of consumers who are shoppers in the way you'd expect.

So either as the fraction of shoppers goes to 1 or as s upper bar goes to 0, prices collapse to Bertrand pricing. So if this is what it looks like with what I've drawn, if I make s upper bar down here, so if I make all the consumer search costs tiny, what you're going to find is that, if this is c, the distribution looks like that.

That is, the distribution collapses towards a point mass on p equals c as the s goes to 0. It also collapses towards a point mass on c as the fraction of shoppers goes to 1.

These equilibria of this kind of model often have U-shaped distributions when N is large. In some sense, if you're a firm and you know you're competing for a large number of consumers, you have two good strategies.

One strategy is typically priced near p lower bar and priced near p lower bar in hope to have the single lowest price and to get all the shoppers to buy from you. And then the other strategy is price up here near the monopoly price and figure that, I'm never going to win all the shoppers. I'm just going to get the people who come to my store first and have a high s.

And so often, you get-- there's often less point of being in the middle, because in the middle, the shoppers-- you're never going to get the shoppers. And the nonshoppers, you would have gotten anyway with a higher price. So you often see these distributions that sort of look like that. They're bimodal.

The limit as N goes to infinity is perhaps not what you might expect, but it is interesting itself. As N goes to infinity, with the way this distribution changes, is it goes like this. You have a small number who go for the nonshoppers. And then you have a huge mass of them up here at the monopoly price.

So this says, in some sense, Tobias talked about cereal product proliferation and there being a million cereals or whatever. This says, if you have a million credit card companies out there, maybe what you would see is almost all of them are ripping people off and just charging the monopoly price. And there are a few of them out there competing for the sophisticated consumers. And we could see this bimodal thing, where most people are at the monopoly price. And in fact, this model does say, as N goes to infinity-- as N goes to infinity, they almost all end up up here.

One argument for that is, if you think about it, let's suppose that you had even a constant fraction of the firms. Suppose that alpha N firms were charging close to the minimum price. Then, if you're a firm and you charge close to the minimum price, you have 1 over alpha N probability of being the lowest firm.

But then because they're also-- because there are, like, alpha times N firms here, there are a lot of firms here. The distance between them is very, very small. So the price sensitivity is very, very large. So this point mass actually has to collapse towards-- this thing becomes very close to c.

And so what you basically get, you get 1 over N consumers, and you get a 1 over N markup on them. And so you'd be getting 1 over N squared profits.

And that's not going to work. Because if you had 1 over N squared profits, you could get 1 over N kind of profits by just pricing up here and getting the people with the high s's. And so it turns out that the equilibrium just is a very, very small point mass here and a very, very big one up there, when N gets big.

So anyway, let's see. I guess I didn't actually say much about the theorem. All I said is that it doesn't have a pure strategy equilibrium. It does have a mixed strategy equilibrium. And in this mixed strategy equilibrium, you get something that looks like-- I just said, it has p lower bar greater than c. I didn't say much else about it. It's got a density.

Why does that happen? The argument for part a, that it's no pure strategy equilibrium, is just like Diamond's. It's like Diamond's, most of the time, that you can't have an equilibrium. Again, suppose we tried to have an equilibrium where someone was at $35, someone was at $36, and then some other people were up here.

You're the single lowest person in this model. You strictly want to increase your price in epsilon overcut. Because, again, if consumers expect this distribution, you're epsilon more than that, you still get all the shoppers, and all the nonshoppers still buy from you.

And so then the only new case is, well, let's suppose we had-- like in Diamond, we have a whole bunch of firms all tied at exactly the same price, whether that's the monopoly price or something else. Well, then you get the Bertrand argument that you would epsilon undercut and get all the shoppers.

So we have the same effect as Diamond, saying, you can't have people different from each other because the lowest firm always moves up. And then you have a new-- and then you go back to the Bertrand argument of, you can't have everyone the same at anything other than c because someone at the bottom goes down. Someone will go down and steal all the shoppers instead of splitting the shoppers with the other firms. So there's not going to be a pure strategy equilibrium here.

What about mixed strategy equilibria? Well, in-- yes.

AUDIENCE: Quick question. How does the shopper's utility work? Because if they get positive utility from opening doors, can't they just do that forever and get arbitrarily large--

GLENN ELLISON: Yeah. So I guess what I have to do is I have to have my shoppers-- you only get s the first time you open a door. You don't-- just-- yeah. Once a door's opened, it's already opened. You don't get new utility from opening it again, I guess. Again, it's a little-- it's mostly a construct to say that the shoppers just open every door.

Sometimes, I guess the way people also think about these models is you may have people who know how to use some search engine that just gives them all the prices instant-- one search gives them all the prices, and that's where the shoppers are supposed to be.

So anyway, in [14.]122, I talk about, this is a game with discontinuous-- profits do move discontinuously in price when you cross where another firm is in a search model. But Dasgupta Maskin's existence theorem for discontinuous games, the sum of the payoffs is continuous. So this is a classic example, where we know from general existence theorems that there must be some equilibrium. It's not pure, therefore it must be mixed.

So we've got, some mixed equilibrium has to exist. If you think about firms mixing over some interval p lower bar, p upper bar, could p lower bar be c? p lower bar can't be c, because if p lower bar were c, firms charging p lower bar would be earning zero profits.

If you're earning zero profits, again, you can go to c plus epsilon. Anyone who visits you first would have to buy at c plus epsilon, so you would get positive profits. So this equilibrium has to have a positive-- it has to be a positive profit equilibrium. So p lower bar is going to have to be strictly bigger than c.

And then why is it a continuous distribution instead of a distribution that has maybe a point mass? There's a large mass of firms who charge here. Well, again, this is like an epsilon undercutting argument. If there was a large point mass of firms who all charged exactly $37, you would go to $36.99. And if in the realization of the randomizations, everyone randomized to $37, then by going to $36.99, you increase the probability that you're the low-priced firm and win the shoppers.

So you can't have a point mass because that would create an incentive to epsilon undercut the point mass. So it's got to be some continuous distribution of prices. Any questions on that?

So what does the density look like? The one case where it's-- the model can be hard to solve. One case where the model is easier to solve is I look at a special case where I just have a two-type version. So I have this two-type version of search costs. I have this fraction mu, have s prime less than 0, and I'm going have the fraction 1 minus mu have some number s that's greater than 0.

So instead of having the heterogeneous s's on a positive interval, let me just make them all exactly the same. I think it's a worse model descriptively, but it's an easier model to solve.

So what happens there? Well, one observation is, what are the profits if you charge price p lower bar? Oh, actually, sorry, first observation. Suppose I go to this. And why do I go to this two-type distribution? So I go to this two-type distribution because, when you have the two-type distribution, every nonshopper, every consumer with a positive s has to buy from the first firm that they visit.

Why do they have to buy from the first firm that they visit? Well, suppose the first firm that they visit has price p upper bar and you ask, OK, I know that this is the worst price that any firm in the world offers, and I just found it. Do I buy from it or do I look for a better price?

If everybody has the same s, the answer is yes, you buy anyway. Because if the answer was no, you don't buy if you find the worst possible price, then the firm who's pricing at p upper bar gets zero customers and is earning zero profits.

So if we're in an equilibrium, it must be that the distribution has shrunk down to the width where if I find-- even if I find p upper bar, it must be that the consumer surplus from finding p lower bar minus the consumer surplus from finding p upper bar is at most s. Because if that wasn't true, no one would buy at p upper bar, and the p upper bar firm would be earning zero profits.

So it must be the distribution just shrinks down. If s is small, maybe it shrinks down to something very tight. But it's got to be that this equilibrium distribution has to shrink in such a way that people always buy when they find the p upper bar price.

And so if everyone buys when they find the p upper bar price, then we have these guys in equilibrium, these search all N. And these guys only search one firm, and they buy if p is less than or equal to p upper bar. So this makes it easier to write down the profit functions because we've gone from a complicated search problem to some consumers are buying from all N firms and some consumers are just looking at one place and buying there, as long as the price isn't above what was expected.

So what do you get as profits if you buy-- if you find the p lower bar-- if you find the p lower bar-- sorry. If you're a firm, you set price p lower bar, what's your profits? It's p lower bar minus c D of p lower bar times the number of consumers you get, which is all the shoppers and a 1 over N share of the nonshoppers. Because if the nonshoppers go to any other firm, they will buy from that firm.

But now in a mixed equilibrium, firms have to be indifferent over all the prices they're mixing over. So that means that at any other price, price minus cost times D of p times the number of consumers that you get, which is you get the people who visit you first and you get the shoppers, mu, if every other firm happens to choose a price above p. And that happens with probability 1 minus capital F of p to the N minus 1 power.

So this is my profit if I set a price of p. This is my profit if I set a price of p lower bar. These have to be equal. And if you look at this equation, the capital F of p appears in exactly one place. So what can I say? First thing I can say is that I divide both sides by p minus c D of p, and I get mu times this equals this over p minus c D of p.

And then I can subtract 1 minus mu over N from both sides. And then I can divide both sides by mu, and then I've got 1 minus f of p to the N minus 1 all alone on one side and stuff on the other. And then I take the N minus first root of both sides, and then I get 1 minus F of p equals something. And so I can solve this for F of p.

So this is an equation we're just F of p appears exactly once. So I get F of p. I can always solve for what F of p must be, given any p lower bar and given any mu and all those other parameters.

And then what is p upper bar going to be? p upper bar is going to have to solve this-- p minus c D of p upper bar equals 1 minus mu over N times this. And so then I can also use that to figure out p upper bar is a point where the revenue is equal to this divided by something or other.

So if p lower bar is too big, that won't exist. But if p lower bar is small, that will exist. And you're going to find that I plug in a p lower bar-- say, I plug in a guess of what p lower bar is. And then given that p lower bar, I say, OK, if that's the p lower bar, this is what the distribution is going to look like. And this is what p upper bar would have to be.

So for any guess of p lower bar, I can solve for what does capital F look like and what does little f look like? And where do we run out of probability? And it has to equal 1 up here.

So is that thing that I solve for-- it's going to be a complicated formula, with something to the 1 over N minus 1 power. Is that an equilibrium? I can always make people indifferent over all prices in the interval p lower bar, p upper bar, but that's not sufficient to show that it's an equilibrium. Because what I have to show is that I've made the firms by construction in different overall prices. In this interval, I have to show that they don't prefer pricing lower or pricing higher.

The constraint that turns out to be hard to satisfy is the price that they don't want to price higher. Why set price p upper bar? Why not go to p upper bar plus epsilon or all the way up to pm?

Because, again, you know that there's zero probability you're going to get the shoppers if you're charging p upper bar. You're only going to get the people who open your door first. Why doesn't it then make sense to open your-- given that you only get the people who open your door first, why not go epsilon higher?

And there could be two reasons why you don't go higher. One could be that p upper bar that you solved for was equal to pm. If the p upper bar was equal to pm, then obviously you don't want to go epsilon higher.

The other thing that could happen is, when consumers see price p upper bar, the gain that they would get from searching again-- again, the gain that they would get from searching again is this. Sorry, it's not that. I will say what it is. But if consumers-- it's going to be the-- it's on the next slide. It's going to be the integral from p lower bar to p upper bar of consumer surplus of p minus consumer surplus at p upper bar F of p dp.

If I go to search again-- I'll call this thing v of p upper bar. If I search again, I may find a lower price than p upper bar. And if I do, I gain more consumer surplus. And so this integral would tell me what is my gain from search.

So it could be that v of p upper bar is exactly equal to s. And if v of p upper bar is equal to s, everybody who finds p upper bar would buy. But if you saw p upper bar plus epsilon, you'd be, that's ridiculous. I'm going to go take a new draw from the price distribution.

So what we have to do, basically, is look for a p lower bar that makes one of two things true. Either it makes p upper bar equal to pm or it makes v of p upper bar equal to s. So anyway, to find the distribution, I just search for a p lower bar that does one of those two things.

This one is actually not hard to find because, if you look at it, we're going to have this equation is true. So I could just plug in the monopoly price here and put monopoly price minus c D of pm times this equals this. And so the only thing that's unknown is the revenue at p lower bar. And so I can solve for the revenue at p lower bar. I can find p lower bar. So anyway, you find an equilibrium by looking for something that solves one of these two extra conditions.

Questions on heterogeneous search models?

Value of Price Search. I put this slide in because this is a "read it yourself at home." But just this is the way you normally think of the value of one price quote. The value of getting a new price quote, if you've already found a price p, is the integral from p lower bar to p, consumer surplus at x minus consumer surplus of p f of x dx. That's the natural representation of what getting one more price quote gets me. It gets me what I gain if I find the lower price quote times the density of the prices I'm drawing from.

Something you see often when you're redoing this, when you're working at papers, is people often write v of p is the integral from p lower bar to p D of x F of x-- capital F of x dx. You can get from here to here by integrating by parts. You can also get from there to there by reversing the order of integration.

People often write down this less-intuitive formula, just because it's well-known, without noting why they're doing it. So I put it on a slide and put demonstration of flipping the order of integration to prove those two things are equivalent.

So second thing I want to do is search with differentiated products. I said that we want search costs because we think differentiation is not sufficient to explain what goes on in the world with real-world markups. But that doesn't mean differentiation is not real.

And so while I think more people do search in that model like I just did, do the Bertrand search and then figure out what it does to markups, it's more realistic to have search added on top of a differentiated product model rather than searching a Bertrand model. So that's what I'm going to show for you here.

So now, I'm going to have two firms with constant marginal cost c. And I'm going to do the Perloff-Salop version of Hotelling's model, where consumers have a two-dimensional type. So consumers have a-- yeah, so they have a theta-- consumer i has a theta i1, and they have a theta i2. The theta i1 is related to how much they like firm 1's product. The theta i2 is related to how much they like firm 2's product.

I guess here, I did it with a v minus t, so this is your-- some people find both firms ideal. They're the ones who have theta is equal to 0. Some consumers really like firm 2, and they don't like firm 1. They're down here. Some consumers really like firm 1, don't like firm 2. They're up here. Some consumers really find both products far from their tastes.

So it's the Perloff-Salop version of Hotelling, where I just have-- instead of having the theta i2 equals negative theta i1 or t minus theta i1, I make independent disutilities-- independent tastes for product 1 and product 2.

So the firms simultaneously choose their prices, p1, p2. And what the consumers here have to do now is, when they pay s, they both learn the price and they learn whether they like the product or not.

So this is the case where you know there are two firms-- you know there are two firms that are selling some product. But you have to actually go to the store and look at the product, at their version of the product, to decide whether you like it or not.

Or you could think of it as, like, in the credit card example. I pay s. When I pay s, I learn-- I learn what the interest rate is. But then I also learn other features of the credit cards that are not-- which people have horizontally differentiated preference.

Like, I do learn what color it is, and I do learn what their prepayment policy is and what day of the week they want their things due. Or I go to the car dealership, and I do learn how hard it is to drive that car. I do learn how clean is the service department at that store and other things or something like that.

So I pay my s, I learn the price, and I learn whether I like it or not. And until I go there, my prior would be that my theta is uniform 0, 1. I know that this might be a good match to me. It might not be a good match to me, but I don't know until I-- I have to inspect the product to see that.

Assume that the search cost is less than t over 2. And again, assume that consumers use some optimal sequential search procedure, which is they go visit one firm, they learn the theta, they then decide, given what I found about the price and the match quality to my tastes, do I go search for a second time, and then they're done.

The theorem is, for large enough v, this model has a symmetric pure strategy Nash equilibrium. And that symmetric Nash equilibrium is price equals c plus 1 over 2 minus square root of 2s over t times t.

So it's a c plus something times t markup. And it's a pure strategy equilibrium. So the Stahl-like dispersion has gone away. And here it's just search cost effect markups. So notice that the Perloff-Salop model with this game would have p star equals c plus 1/2 t.

Here, if you set s equals 0, you get c plus 1/2 t. So when s goes to 0, this goes to Perloff-Salop. But when s is larger, like, imagine s is t over 2-- when s is t over 2, you get 2 times t over 2 is t over t. So it goes to c plus t. So price increases smoothly from c plus t over 2 to c plus t as search costs increase.

This model has the "falling off a cliff" property that, if you make s-- so as s goes from-- if I graph a search cost, when s is 0, price is c plus t over 2. As s goes up from 0 to t over 2, it's just a 45-degree line. Prices increase and they reach c plus t.

Once I pass this search cost, we fall off a cliff and the search disappears and no one ever buys. And the reason is because, when s is large enough, no one is going to get a second quote.

So even if you find the first product and it's the worst possible product for you, and you know, OK, I hate firm 1's product, the most you're going to gain from getting firm 2's product in expectation is t over 2 because firm 2 is disutility-- firm 2's product might be just as bad. Firm 2's product might be great. But you're only going to gain t over 2 on average.

If you're only going to gain t over 2 on average, you won't do a second search if your s is bigger than t over 2. And so then once people have an s bigger than that, we're back to Diamond. Everyone only searches once. We go to the monopoly price.

But when we get to the monopoly price without the multiunit demands, at the monopoly price, you don't want to be-- there's no point in buying at the monopoly price because the monopoly price is at least v minus t. Your expected utility's, at most, t over 2. In this model, at the monopoly price, you wouldn't bother searching to begin with. Because consumer surplus, the monopoly price is very low in this single-unit demand model.

So intuition for why is it that equilibrium prices increase in search costs? If a firm is considering a price increase, they have to trade off the gains that they get from these inframarginal consumers who would have bought anyway against the losses from raising the price and having the marginal consumers go away.

Here, as s gets bigger, there are fewer marginal consumers. When s is small, most consumers are going to visit both-- when s is small, I'm going to have some region like this, like, these consumers buy and these search again.

Because if I'm firm 1, for someone to buy from me when s is tiny, why not get a second price quote? It's only if they really like me that they'll stop and buy from me. When s gets big, this line shifts out, and it becomes these guys all buy. And it's only if you have a really bad draw that you search again.

And so what happens is just, when s gets big, there are fewer and fewer consumers who see both prices. Because there are fewer consumers who see both prices, there's a smaller gain to an epsilon price cut. And so your incentive to cut price is smaller because there are fewer consumers who actually see both prices. Whereas, you know in equilibrium, you're getting half the consumers. And on those half the consumers, if you raise your price, you're going to get more on all-- on roughly half of the people.

So the number of marginal consumers goes down. The number of consumers you're getting in equilibrium stays the same. Therefore, your profit on the marginal consumers must be going up to make you not want to lose them. And so that's an intuition for why the markups have to go up because you-- the markups have to go up to the point where you don't want to lose that dwindling number of marginal consumers. Any questions on that?

I'm trying to think. Why don't I go over this in a little bit of detail, just how the calculation in a model like this works. So imagine in this model you've got a consumer. I switched from the theta dimension to the match quality dimension. So imagine that a consumer visits one store. They find a match quality of x, by which I mean 1 minus theta i, and a price of p star, which is the price they expected to get. What's your expected gain from search?

In equilibrium, you expect that the other firm is also charging p star. So the gain comes from getting a better match quality at the second firm, which is just an integral from x up to 1 of u minus x du.

So I might find a firm that's essentially no better. I might find a firm that's 1 minus x better. And I'm equally likely to find a firm that's any-- well, I may find a firm that's worse than what I found already. But if I found a firm that's worse, I can just ignore it.

If I find something that's the same or better, I can buy from it. And so there's a 1 minus x chance that I'll find something better. The most by which I might increase my utility is 1 minus x times t. And so I get this integral, which is just-- it's the area of a triangle, 1/2 t 1 minus x squared. So that's my gain from search is the gain from improved match quality.

When are you going to do a second search? Well, you're going to do the second search if-- suppose I find a match quality of x, and I actually find a price that's different from p1 star. So if the firm is charging even a higher price, it's like the match quality is worse. If the firm is charging a lower price, it's like the match quality that I found is better. And so I'm going to buy-- I got to search again if v of x plus p star minus p1 is less than or equal to my search cost.

In equilibrium, p1 equals p star. So this holds if 1/2 t 1 minus x star is less than s, which is if and only if x is at least 1 minus square root of 2s over t. I just solve this thing for x, moving the 2 divided by t over to the other side and then taking square roots and then doing 1 minus it.

So the Nash equilibrium is always the solution to partial pi 1 with respect to p1, evaluated at both firms charging p star equals 0. And as I said, often, if you're looking for symmetric equilibria, rather than writing out best responses and intersecting them and solving two equations and two unknowns, you just take advantage of the fact that you know it's symmetric and just do, this first-order condition characterizes what the symmetric equilibrium p star must be.

And so let's think about, well, if p star is in equilibrium, what happens if firm 1 cuts its price to p star minus dp? And let's think about, what are the gains and losses when I'm comparing profit at-- cutting my price to profit from keeping my price the same?

And as I said, one loss that I get from a price cut is, I know in equilibrium, half the consumers are buying from me. So if I cut my price by dp, I'm giving dp away to all those people. That's a loss of 1/2 dp, because I'm getting half the consumers in equilibrium. Again, exploiting symmetry.

What do I gain? OK. So I've got three different gains. First gain is, suppose that a consumer visits me first. That's half the consumers visit me first. When consumers visit me first in x space-- in x space, it looks like this. There's some x star. And here they buy from me, and here they search again. If my match quality is good, they buy. If my match quality is bad, they search again, where x is 1 minus theta.

So match quality x. High match quality is you buy. Low match quality is you search again. So where I gain from the price cut is they visit me first with probability x star. They didn't like my product all that much. So they continue to search.

But then it turns out that the amount that they like the other firm's product is just ever so slightly greater than x star and so close to x star that my dp price cut convinces them to buy from me instead of him. And that's going to happen with probability dp divided by t. So it's going to take this coincidence of things that they visit me, they then visit him. And then they like him just about the same, so my dp price cut makes a difference.

There's also consumers who visit-- oh, sorry, I did that one backwards. So there are also consumers who visit the other firm first, don't like the other firm's product, visit you, like your product again, almost exactly as much as they like the other firm's product, and then the dp price cut makes a difference. So we get these two terms that are both 1/2 x star 1 over t dp.

The third effect is what we'll call the search deterrence effect. The search deterrence effect only comes from consumers who visit me first. And they're consumers who visit me first and would have bought from firm 2 if they went and visited firm 2 but who I convince not to visit firm 2 with my price cut.

And so when I cut my price by dp, it's like I've made the match slightly better by some amount. It's like I increased the match quality to x star plus dp over t. So I've made my match slightly better. I've shifted the-- oops, I've made my price slightly better.

So by making my price slightly better, I'm shifting this region back. So if the-- yeah, if my price is-- if my price is lower, the match quality in the-- quality space, they need to-- it's, like, x star shift back. Because if my price is better, they'll take it even if it's a slightly worse match quality in characteristic space.

And so I shift that line to the left by this distance. And then if they would have bought from me anyway, it doesn't matter that I deterred the search. But if they would have bought from the other guy, then it is a gain.

And so I get these two effects. One is the standard Bertrand effect-- or the standard Hotelling effect of the people who see both prices and they're close together by the epsilon price cut. I get those consumers. And then I have this extra effect of search deterrence effect. That is, by cutting my price, I keep people from searching the second firm. And so it must be that these three gains exactly offset this one loss.

Anyway, and so what you do is, the first-order condition sets the gain equal to the loss, you solve, and it gives you that formula. So that's the mechanics-- that's the mechanics of finding an equilibrium. And I guess the key economic point is just there are those two tradeoffs-- the price comparison effect and the search deterrence effect.

Fixed Sample Search. So anyway, any questions on search with differentiated products?

So fixed sample search, in some applications, the model we've been doing so far has this idea of, you go out and you do a search. You see what you find, whether you like the product. You decide whether to stop here and buy or go do another search.

There are some applications, like getting contractors to renovate your bathroom, that you can't do this sequential search process. It takes a long time to get a contractor to submit a bid. You call them on the phone. It takes the guy forever to come. He looks at your bathroom. He says, I'll get back to you. He doesn't get back to you.

A month later, he finally works out what he's going to charge you. If you do that sequentially, you know your bathroom will never get renovated. So what you do in practice-- and as governments-- we do this all the time. We're renovating new space in the economics department. You try to get two or three contractors to bid.

You say, OK, we're getting three bids. You may put a bid deadline on them. You may tell the guys, look, I'm getting three bids. I want all the bids in by December 1. If you don't get your thing in by December 1, I'm not going to consider you. By December 1, you get three bids and then you compare them and then you're done.

But you're deciding on how many bids you're going to get before you see what the bids are. And that's the difference between what we call the fixed sample search model versus a sequential model. Sequential model, you treat them one at a time and decide whether to go on. The fixed sample one, you decide how many to get.

And it turns out that these models work somewhat similarly. For a while, this fixed sample model was more popular in structural empirical work because the equations turn out to be a little bit-- if you think about it, in the sequential model, the consumer is doing this dynamic programming problem. In the fixed sample model, the consumer is choosing a finite integer.

And so the choosing a finite integer is a much simpler strategy space for the consumer than this optimal dynamic programming thing. And so therefore, there are a lot of models that use the fixed sample search. Sometimes, it's appropriate. Sometimes, you think sequential seems better. But fixed does a lot of this. It behaves in a lot of ways similarly.

So let me go over how you would solve that one. So again, the game is going to be-- oh, actually, timing of the game. I should have reversed this. First stage, consumers decide how many searches k they're going-- quotes they're going to get. Oh, no, actually, timing the game first.

So first, the firms simultaneously choose prices. So I did the timing right. The firms are choosing prices, and they're selling a homogeneous good-- renovating your bathroom. Let's say it's homogeneous.

You have heterogeneous consumers. They have search costs s that are distributed according to G on some interval s lower bar, s upper bar, like in Stahl. No shoppers here. And then the consumers are going to, again, have multiunit demands. And what the consumers do is they choose how many price quotes to get.

So the timing of the game is, the firms choose their prices. Each consumer decides how many quotes they want. They then get the quotes by sampling randomly from the firms-- set of firms. And then the consumers buy from whichever firm gave them the best price.

So again, as in that Stahl model I did, this is going to be a model that has a mixed strategy equilibrium. And it's going to have a mixed strategy equilibrium-- it doesn't look exactly like these, but it's going to be the same kind of thing. There's going to be some density f of p on some interval p lower bar, p upper bar.

If that's what happens, firms mix over some interval p lower bar, p upper bar, what's your value from getting k price quotes? Your value from getting k price quotes is you're just going to integrate over from p lower bar to p upper bar, consumer surplus of x times f of 1 colon k of x dx, where f of 1 colon k I'm defining to be the first-order statistic of k draws from that distribution.

So if I'm getting k draws from that distribution f, what is the density of the lowest of the k draws? And the density of the lowest of k draws from that distribution is-- what's this? Probability that the lowest draw is exactly x is k times f of x times 1 minus f to the k minus 1. So there's one firm offering price exactly x, which happens with probability k f of x, and then all other firms have higher prices.

So if this is the distribution, and then I say, what's the distribution of the best of two draws in that distribution, it's got more mass here and less mass here because I'm more likely to find something big. If I do a hundred draws, I'm very likely to find the lowest price and very unlikely to find the highest price. This is the best in my sample of a hundred.

So anyway, I'm choosing k. That affects the density of the lowest price. And then this becomes my utility that I'm going to get from buying from that firm, I guess, minus the search cost.

Notice that this model has a natural decreasing marginal returns to search property. The value of the first search is bigger than the incremental value of the second search is bigger than the incremental value of the third search and so on.

Why is that? Well, it's because the incremental value of the k-th search is just I'm integrating over all prices that I might find on the k-th search v of x times the density of the best price. So think about, I'm just conditioning on the best price I've found in the first k minus 1 searches. And so the best price I found in the first k minus 1 searches was x. I'm now going to get v of x from an extra search, where v of x is that quantity I gave you before, that integral.

And so as v is an increasing function, the higher the best price I've found, the higher is my gain to search. But then as k increases, these distributions are decreasing in a first-order stochastic dominance sense. And so because I've already found a good price, the gain from a new search is reduced.

So once I've already found-- if I've already done a thousand searches, I've probably already found something very close to the lowest possible. So the value of one extra search is small. And it just has that-- it's got that monotonicity property. Every extra search, it's more likely that the best price I've found is low. Therefore, it's more likely my gain from search is low.

So what do we get in equilibrium? In equilibrium, we get a graph that looks like this, where, if your search cost is less than-- this line here would be at U of N minus U of N minus 1, or U sub N minus U sub N minus 1-- if your search cost is below that, you'll visit every firm. If your search cost is above that but less than U of N minus 1 minus U of N minus 2, you'll visit N minus 1 firms.

If your search cost is higher, you'll visit N minus 2 firms, so on. You go all the way out. There's some consumers who visit one firm, and there may be some consumers with s is high enough that they would visit zero firms.

So what we're going to get is the search costs are just going to divide up in these intervals. So there's going to be an interval of people who do N, N minus 1, N minus 2, and so on. And then remember, I had some density capital G on the search costs on some interval. So then integrating that density over these regions is going to give you the-- so if I-- don't want to write on my screen.

I have these regions where you do N searches, N minus 1 searches, N minus 2 searches, all the way out to this region where you do one search. And then there's some density of search costs that looks like that. This is the g of s. The integral of g over this region is the mass qN of people who do exactly N searches.

And the integral over this region, this is qN minus 1, the mass of people who do N minus 1 searches. And then I go all the way out to here. The integral over this region-- these are the people who do one search.

So just integrating the density g over these boundary regions of how many searches you do gives me the number of people who do one search, the number of people who do two searches, the number of people who do N minus 1 searches, the number of people who do N searches.

And then just like in Stahl, we're going to get-- the fact that-- if you think about Stahl, Stahl was-- the solvable equation of Stahl I did for you was q1 was mu, qN was 1 minus mu, all these other were 0. That's what I did for you, is there are a bunch of people doing one search, a bunch of people doing N searches, and then nobody else, in that one case I solved for you.

In this model, all these q's are going to be non-zero, but you're going to solve this model in exactly the same way. You're going to say, I know what the distribution of-- I know what the profits are. In this case, the profits that I know is I know that the profits, if I charge the monopoly price, are going to be pm minus c D of pm times q1 times 1 over N.

If I charge-- the p upper bar is going to be pm in this model, always. So if I charge the highest possible price, my profit is-- I'm going to sell to the consumer if and only if they visit me first and they're one of the consumers who decided to only get one price quote. So the high s people here decide to get exactly one price quote.

So I'm going to sell to them if they get exactly one price quote and that one price quote is me. And then I'm going to rip them off with a monopoly price. And this is going to be my profit, pm minus c D of pm.

What do I get if I charge any other price? Well, I'm going to get, again, p minus c D of pm per consumer. But then how many consumers am I going to get? Well, if a consumer visits me first-- if a consumer is doing one search and visits me first, they always buy.

But then also, there are some consumers who visit two firms. They visit me with probability 2 over N, and then I sell to them with probability 1 minus capital F of p. Because it has to be they visit two firms-- they visit me and somebody else charging an even higher price.

And then there are also consumers who visit three firms. They visit me with probability 3 over N. And they buy from me with probability 1 minus f of p squared. Because it has to be, they visit three firms, I'm included. And the other two firms turn out to have higher prices than me.

Well, again, this becomes an n-th degree polynomial instead of something that just has 1 minus f in one place. But you can still-- this equals this lets you solve for f of p for every p less than pm. And so it's the same kind of argument as in Stahl. It's like, if you knew q1 through qN, you could solve for what capital F would have to be.

But then in this model, there is a fixed-point calculation. You say, OK, here's my guess for what q1 through qN are. I guess q1 through qN. I figure out what capital F would have to be. Once I figured out what capital F would be, I do this computation of what the U's are to figure out how many consumers search N things. And I see if I reproduce the q1 through qN.

And so numerically, I solve this by finding a fixed-point problem. I guess q1 through qN. I figure out what f would have to be. I figure out what U would then be. I figure out what the q's would then be, and I see if they match. If not, I have an optimizer moving the q's around to try to get the q's to match the q's I started with until I find a fixed point.

Again, it's not something theorists do often because it doesn't have a nice closed form. But it is something that empirical people will often do if they're solving their model numerically anyway. It's easier than the dynamic programming thing that the other one has in it. Any questions?

I have just two brief slides I'll do. and then the other one, I'm actually-- I think I'm going to let [? Rowey ?] talk about on Friday.

So obfuscation. So an observation about all these search cost models is-- so far today, I've always taken the s as an exogenous parameter. And then profits of the firms have gone up in s.

So that raised the question of, where do s's come from? Why are there search costs in models? And clearly, firms collectively are better off when s is bigger. The bigger is s, the higher are the firm profits.

And so certainly, if all the firms could get together and agree, are we going to make searching for a mattress easier or are we going to make searching for a price on a mattress hard, all the firms will get together and say, let's make searching really hard for the consumers and we can all charge a monopoly price and be best off.

But that raises the question of, does that require collusion among the firms to get s high? Or can firms also individually have an incentive to raise their s, even if it's just individually, do they want to raise s?

And then you might think, well, that's going to be harder to get to work. Because if I raise my s and consumers see that I'm the annoying firm to visit, aren't they all going to go to the nonannoying firm first and pay less to get s there? And isn't that bad for me if I raise my search cost and I drive consumers to the other firms?

And there are a few papers that discuss reasons why you actually can get it to be individually rational for firms to raise search costs and not just collectively rational.

So one paper I put is this paper by Chris Wilson in the*IJIO.* And this considers this game, which has a surprising result. Firms simultaneously choose search costs, s1, s2, from s lower bar, s upper bar. The consumers and the firms observe what these search costs are, and then the firms choose an unobserved price, and then consumers follow a search and purchase policy with optimal sequential search.

And it turns out that this model does not have an equilibrium where both firms set price s lower bar. So it is true that you try to have a firm set price s lower bar. You have this one firm, what it can do is it can raise its price. It can raise its price and be known as the annoying firm, so that all the consumers are going to go to firm 1 first.

So if I'm firm 2, I raise my search cost, and then I get everyone to visit the other firm first. But then the other firm knows that it's getting visited first. And so because it knows that it's getting visited first, consumers are going to pay a search cost to go to the second firm. It's going to have a big incentive to raise its price because it knows, I can at least go s bigger than the other guy, and people are going to not go visit him.

But then, of course, once it goes s bigger than the other guy, it's going to have an incentive to raise its price. Then firm 2 is also going to say, well, I'm only getting the people who visited firm 1 and then found out that firm 1's price was unexpectedly high. And so then they're going to raise their price.

And you can create this asymmetric Bertrand-like competition, where the firms have an incentive to raise their price. And then because I know my rival's going to raise their price, I may raise my price too. And therefore, even though they're starting with him, he chooses a distribution-- a mixed strategy that's so high that people come to me anyway, and we're both better off. So that kind of story can work.

The other story that I have in a paper with Alex Wolitzky is that, you can also do this when the cost of getting a price quote is not observable in advance. Suppose that it's only when you get to the car dealer that you learn whether it takes one hour to get a price quote, or five hours to get a price quote from that dealer, and how long they're going to make you haggle before they get the real price they're willing to offer.

And we just discuss in our paper other ways you can do it. One way you can do this is with the convex disutility of time spent shopping, where after you've spent 5 hours in a car dealer, you now really don't want to spend 10 hours by going to the second car dealer. And there's a convex time cost. The other is if you add uncertainty over how efficient it could be to buy a car. And this one I think works well with getting a quote from a contractor.

You get one contractor, and you spent 14 hours with this contractor trying to get a quote from him. You think getting a quote from another guy is also going to take you 14 hours because it must be that hard. And if consumers don't know how hard getting price quotes are, by making it more annoying, people are going to infer that the other dealers are also more annoying. And therefore, that's going to make you also individually raise search costs.

Anyway, so that's the argument there for there are reasons why, even though you lose some customers, you may want to raise search costs. So let me say, plan for the rest of the week.

I think I'm going to have [? Rowey ?] talk about multiproduct search. We don't have any problem set to go over on Friday. So I think he's going to talk about a couple recent papers. This paper on multiproduct search is likely to be one of them. And then next week, I'm going to jump ahead to empirical search. And these are the three papers I plan to talk about.

There's no class Monday due to the holidays, so it's next Wednesday. And then if you want to read one in advance, I would read my paper with Sara on search and obfuscation.