[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK, as I said, so, classic example here, American Repertory Theater in Cambridge, plays that are going to go to Broadway, they try them out there, also sort of avant garde stuff. If I buy a ticket there, it's like $72 to sit in the back corner of the theater. You, as students, they actually, if you go on their website, they have a tremendous number of students, basically everyone except me. And the student discount there, you can go see plays that cost me $72 and you pay, like, $10.

It's less than going to a movie. I would say, it's a great thing to see, that it is typical, though, that you might buy a ticket to see Hamlet and find that you've got guys standing in kiddie bathtubs, wearing shorts, playing guitars, or whatever. So it is-- you have to be aware that you may not get what you think you're getting when you see the title of the play, but a great thing to go. But anyway, example of this, they're selling the same ticket to you for $10 they would charge me $72 for.

People often also think about this when you have nonlinear prices. You'll go to the supermarket. And it's, like, buy one, get one half off. It's like, the first good is being sold at a different price from the second good. And again, the first thing of strawberries you buy and the second thing of strawberries you buy, they're identical, except they're selling them to you at different prices.

We also talk about price discrimination when firms are selling similar goods at very different prices. So, for instance, I put the iPhone 14 down there. The iPhone 14, 128-gigabyte version is $799. The 512-gigabyte version is $1,099. So it's $300 extra for the more expensive iPhone. You might think to yourself, wait, I can buy a 512-gigabyte flash drive for $50. Why does it cost $300 more to put it in the phone? Is it better memory? No, it's exactly the same memory that's in the $50 flash drive that someone's making a profit on. But somehow, two similar goods are just being sold at very different markups.

Again, whether we think of something as price discrimination when the goods differ, do you talk about discrimination if it's, like, the ratio of the prices to costs are different, or if the dollar value of price to cost is different? But it is-- often in these examples, it's very, very clear that, however you think about adjusting for the cost, something is just much more expensive than the others. We'll call that price discrimination. The business class seat is $12,000 when the coach seat is $600 or whatever. It's like, I don't know whether I'm counting cubic inches or whatever. But one is just much more expensive.

Also, sometimes, you can think about it as being price discrimination when costs differ, and you don't discriminate in prices. So for instance, there are places that offer free delivery. And the free delivery doesn't depend on how far it is to get to you. So you might think that they're, in some sense, when the firm is incurring a large delivery cost to get it to you, they're giving you a discount, relative to people for whom they're spending a small amount of money to get it to you.

And so, again, I would think of that as charging the same price can actually be price discrimination, if the costs are very different of serving the different consumers. OK, obviously, price discrimination requires market power. This was actually one of my favorite signs I found at the Whole Foods once, which was cranberries for sale, regularly $2, now 2 for $4, which, I didn't quite know what was going on there.

OK, so, theory of price discrimination goes back to Pigou in the 1920s. Pigou made up three names for three different types of price discrimination. These names are terrible names. But they've survived. So, first one, he called first degree price discrimination, which some people also refer to as perfect price discrimination. This is the ideal of a theorist where the firm knows everything they need to know. They have all the power they would want. They can price separately to every consumer. It's really a theoretical ideal because firms are never in the situation of exactly knowing the consumer's utility function.

Second degree price discrimination is price discrimination based on self-selection. So this would be offering an iPhone 14 in two different quality levels and relying on the consumers to sort themselves between, do they buy the cheap iPhone 14 or the expensive iPhone 14? And you see lots of other examples of this, cell phone plans, health insurance contracts, when you buy regular premium gasoline. Firms are setting these prices, relying entirely on the consumers to self-select among them.

And then third degree price discrimination is the examples like the ART student tickets or military discount or senior citizen discount or whatever, where there are some observable characteristics of the consumers. And then the firms can set different prices based on the different characteristics of the consumers. Another one you'd see here, classic third degree discrimination, is geographic location. Firms will set different prices in different countries because they can observe in which country someone is buying the good.

I would say, why is this a bad nomenclature? One is, names first, second, and third degree don't really have any meaning. And then, besides not really having any meaning, the order is wrong. In some sense, first degree is when you know everything about the consumer. Third degree is when you know something about the consumers.

Second degree is when you know nothing about the consumer. So they really ought to be in this order. But anyway, that's the order Pigou described them in. People still talk about second degree price discrimination. So, anyway, I'm going to continue to use those names because that's what everyone else uses.

So let me start with first degree price discrimination. So, first degree is the simplest. So suppose this is the omniscient monopolist, the monopolist, powerful, whatever. So, a monopolist, every different consumer who walks into the store, the monopolist knows, this is p1 of x, this is p2 of x. So if customer 1 walks into the store, the monopolist knows, OK, I'm facing customer 1. This is customer 1's willingness to pay for x units, for any number, x.

And then, customer 2 walks in the store. This is customer 2's willingness to pay for any number of units, x. So I'm going to suppose that this monopolist is powerful. It can set customer-specific nonlinear prices. And it can prevent resale across the consumers. So what's going to happen? You can imagine some very complicated menu of prices you set to consumers. You have them play some game. You have them do something or other.

But at the end of the day, suppose, in equilibrium, this consumer buys x1 units and ends up paying x2-- sorry, ends up paying t1. And this consumer buys x2 units and ends up paying t2. If this is what's going to happen in equilibrium under some very complicated scheme, you might as well just make a direct take it or leave it offer to each consumer.

So just offer consumer 1, either you buy x1 units at a price t1 or you get nothing at all. You go to consumer 2. Either you get x2 units at a price of t2 or you get nothing at all. And so, anything you can achieve with a more complicated scheme, you can achieve with these simple take it or leave it, customer by customer take it or leave it schemes, OK? So that's what we're going to think about doing, is just offering these customer-specific take it or leave it schemes.

When is customer 1 going to accept an offer that involves x1 units at a price of t1? Well, if that's my offer, this is the maximum amount that the customer is willing to pay for x1 units. They're willing to pay this for the first unit, this for the second unit, this for the third unit, and so on. So the maximum they'll pay for x1 units is just the integral from 0 to x of Pi of s(ds).

Given that the customer is going to accept that offer, if and only if Ti is less than this, clearly, the best thing to do is to set the price as high as the consumer's willingness to pay. So you set Ti equals the integral from 0 to x, Pi of s(ds). And so, thinking about the monopoly problem, you're just trying to maximize the total amount. You pick the number of units, xi, to sell that consumer. This is the price you're going to charge when you sell x units. And then this is the cost of producing the units.

So, clearly, I differentiate that with respect to xi. The derivative of this with respect to xi is just Pi of xi. The derivative of this with respect to xi is c. So the first order condition is just, set xi-star, such that Pi of xi-star equals c. So in some sense, it becomes, the optimal solution is, do exactly what would happen with perfect competition.

So if this is my constant marginal cost, for every consumer, I just figure out, how many units would they buy if I was setting the price then equal to c. And I do this as x1-star. I do this as x2-star. And then the total amount that I charge them is just their total utility they get from buying the number of units that I sell them.

So, comments, quantities are exactly the same as with perfect competition. First degree discrimination is socially optimal. Again, I'm ignoring distributional effects. But it's just, there's no failure to realize gains from trade. Every unit that is worth c gets sold to somebody. Notice also that, rather than having this take it or leave it offer bundle, the monopolist could also do this by just offering the good at a price of c, and then having a fixed fee that they add on top of that.

So in this case, for this consumer, what I can do is say, you have to pay me this fixed fee, which is equal to the area of this region. And then you can buy as many units as you want at a cost of c. And so, for this consumer, again, you have to pay me this fixed fee. And then you can buy as many units as you want at a cost of c. So for instance, suppose you're a cloud computing provider. And you have customers who are going to be doing a lot of calculations on your platform.

What you just have to do is, you want to charge all the cores that they rent from you and all the memory they rent from you at your marginal cost of providing cores and rent to avoid distorting what they do and make the service as valuable as possible for them. And then you make the service as valuable as possible. And then you extract all of the surplus they would have gotten, using the service in a fixed fee up front.

So it costs you $2,743,000 to join my platform. And then you can rent computing costs at a marginal cost. It costs you $4.7 million. You just have a different price for every consumer to join. And then, once they're there, you set the price equal to c.

In some sense, you could think of that as maybe that's what happens at Disneyland. Disneyland, you pay $135 to get into the park. And then the rides are free. And the rides are free because the marginal cost of providing a ride is essentially 0. So what you want to do is make the rides free, make the park as valuable as possible, and then extract it all with these fixed fees, OK? The difference between Disneyland and this example is, Disneyland tends to charge everyone the same price.

Here, you would have a customer-specific price. Every family that walks in, you'd be like, oh, for you, $833 to go to the park, for you, $445 to get into the park. And every customer would have a different price. But other than that, the rides are free. OK, questions on first degree discrimination?

In reality, there are a number of reasons why we don't do this. One is, at least before there were these omniscient internet firms, consumer preferences are typically not observable. We don't know what the goods are worth to the consumers. Therefore, we can't do these individual, personalized prices. Another constraint is arbitrage or resale across consumers. So, imagine you did-- Netflix is $12.95 a month. And you get to watch an unlimited number of movies.

You have the problem that, instead of the one person watching an unlimited number of movies for $12.95 a month, all their friends and their relatives can all use the same password and all log in to Netflix and all watch an unlimited number of movies per month. So you need to be able to prevent people from reselling the good across consumers and having different people consuming off the same pattern.

You could also imagine, you might want to have something like-- it seems like, with these statistical packages, I can get this much power for a very low price. And then, for a low price, you can get a low number of units. And then, for a high price, you can get many more units. You don't want someone just coming up to you twice and buying the good, and then showing up again half an hour later and say, OK, now, I'd like more units also at the low price.

And then administrative costs. If you think about, where does price discrimination occur, in the United States, many, many goods are sold in stores that have prices printed on them. And you can walk in. And you buy the good. And you go home with it. You go to many developing countries, that's not the way stores work. You walk into a store, or a guy's shop. And you want to buy something, especially if you're a tourist, and they just start discussing what the price will be. And there is a long haggling time. And then, at the end of haggling, you eventually get to buy what you want to buy.

Why does that happen? I would say, probably, it's the low cost of labor in developing country retail. There's just tremendous excess employees there. You have these excess employees whose time is worth almost 0. They can do the hard work of trying to price discriminate, figure out what it's worth to you, make some extra money. In the US, we actually we do this for-- we do haggle over cars and over houses and things, where it's like, it's a car. The car is $30,000.

That's enough of a difference that it makes sense to have a salesperson who's trained in negotiating and figuring out how much the customer's willingness to pay and trying to get them to pay a little bit more, or pay them a little less, or get them to pay a little less if they need to. But normally, most goods, it's not worth it. I think there's a thought that price discrimination may have gone up with the rise of the internet because a bot that prices things in a discriminatory way doesn't charge by the hour for its work, trying to figure out how to price discriminate among people.

Obviously, you can come up with several examples where, if these problems are severe, price discrimination is impossible. So, for instance, simplest example here, suppose customers have unit demands. And there's a competitive resale market. If there's a competitive resale market, then in that resale market, the price of the good is going to end up being capital P, where P is the inverse demand for aggregate demand. It's going to be P of the total number of units that got sold in equilibrium.

And so, given that consumers can always buy this good at P* in the resale market, they're going to buy it from the monopolist, if and only if the number of units the monopolist is selling them is less than P* times the number of units they get, if that's the price. Therefore, you might as well just directly sell it for P*. And there's nothing you can do. Anyway, I put a couple other examples on there.

Here, I always like finding examples of factors limiting price discrimination and how, sometimes, it's harder to price discriminate than you think. And sometimes, it's easier. So here's the, it's harder than you think. There's a story of an Italian supermarket that had a 20% discount on groceries for anyone over 60. And they thought that this was going to be a fantastic plan because senior citizens are more price-sensitive. They're willing to shop around more. They may go to multiple stores. We'll just offer them all a discount. And they will buy the goods.

And what they realized, apparently only after a year, was that a number of seniors in that town had started businesses of shopping for other people. And they had become professional shoppers. You would go hand them your grocery list. They would walk into the store, buy all the goods, get the 20% discount, and then split the money with the customers, who were happy to-- in some sense, it was a forerunner to Instacart or whatever. They would take the discount. They would give you the stuff. And then everyone was everyone was better off. And I guess, when the firm realized that, they decided, maybe this wasn't such a good price discrimination scheme after all.

Second example, this is a classic photo. It doesn't come up so well. At least, if I turn off these lights, you'll see it better. But for those of you who are familiar with the Dutch and French flags, the Dutch and French flags are very similar, although one is, the stripes are vertical. And the other, the stripes are horizontal. But if you're a company that decides to sell pillows online and you want to make a Dutch flag pillow, it would look kind of like this, a red stripe, a white stripe, and a blue stripe.

If you want to make a French flag pillow, it would look like a blue stripe, a white stripe, and a red stripe. Obviously, at some point, they realized, well, it's the same pillow, just turned on its side. But somehow, on their website, the website had the two of them there. I don't know if Dutch people pay less for pillows than French people do. But they had, the Dutch pillow was $9.99. The French pillow was $15.99.

And I guess you had to count on people not realizing they could buy the wrong pillow and just turn it sideways to get the pillow that they-- the French people could get the Dutch pillow and turn it on its side and get it more cheaply. I mean, the reds and the blues are not exactly the same on the three pillows, on the flags themselves. But I'm sure, if you bought those pillows, they're the same pillow, just turned sideways. And apparently, arbitrage was not enough to prevent that scheme from working to some degree.

OK, so, next step, third degree discrimination. So, suppose the monopolist now can distinguish between classes of consumers but is limited to using simple linear pricing within each group. So for instance, maybe this pillow retailer actually has two different websites. They have a Dutch website that's only in Dutch and a French website that's only in French. And so they can just set different prices in different countries. And consumers can't see the other website.

So anyway, so I'm going to have two groups. i equals 1 and 2. This will also be students and non-students. Each group has a demand, Xi, of p. And then, again, I'm going to go with the constant marginal cost of c. Here, I think what I'm going to do is think about group 1 as being the students, or the low demand consumers, and group 2 as being the higher demand consumers.

OK, so this is population 2. I put price in this axis. I guess I switched. So this is x1 of p. This is x2 of p. So, at any price, the students are willing to pay less than the non-students because, presumably, they have less money. So anyway, with discrimination, what are you going to do? Within each population, i, you're just going to maximize pi minus c times Xi of pi. This is just the monopoly pricing problem.

So the monopoly pricing problem has the marginal revenue equals marginal cost first order condition. Price minus cost over price is 1 over the negative 1 over the elasticity of demand, just like it was last week. And so, let me just say here, let's suppose we have cost here. So this is p1*. And then, this is cost. This is p2*. So I'm going to assume that-- label the markets, that the monopoly price in market 1 is smaller than the monopoly price in market 2.

So third degree discrimination, if you're allowed to practice it, is just monopoly pricing market by market. Interesting questions that come up in third degree discrimination is, what if you decide that discrimination should not be allowed and we want to ban price discrimination? You know what we say? And generally, we have a funny relationship with price discrimination in the US. So there is something called the Robinson-Patman Act, which dates to the early 20th century.

Robinson-Patman Act bans price discrimination. But it bans price discrimination to the extent that it harms competition. And so, it was originally developed for things like not allowing the railroads to charge one firm more for shipping oil than another firm for shipping oil in order to create an oil monopoly on the East Coast of the US, or something like that. It's typically not applied to retail price discrimination or price discrimination that isn't going to be harming competition of something, where you're providing an intermediate input to some downstream firms.

But there's some times when we think that price discrimination is OK. Somehow, it's OK to give senior citizen discounts. It's OK to give student discounts. But it's not OK to charge people extra. So I don't quite know how it is. It's like, discounts are OK. Charging someone more would be totally inappropriate. If you had a women's surcharge that is, we have a price in our restaurant. And women have to pay 10% more. That would be considered completely inappropriate.

Or if there's a big storm, we have a surcharge on gasoline or firewood or whatever, because there's been some natural disaster. Discrimination there is just completely inappropriate. And there are anti-price-gouging laws that prevent that. So anyway, what are we going to do in this market if you decide to ban price discrimination? Well, you're just going to maximize over a single price, p, p minus c, x1 of p, plus p minus cX2 of p.

If you want to think about what the first order condition looks like, you just have-- the single good case we had this was d-pi i-dp You just are adding that up across the two markets and getting the derivative of profit with respect to p in market 1, plus d(pi)/d(p) in market 2 equals 0. What's that going to do? The most standard result is that, if demand is well-behaved, by which I mean two things, demand is single-peaked and it's concave in p when p is less than p2*, then what the firm does is, when you ban discrimination, they just choose an intermediate price. The consumers, the students, are hurt by the price going up. The non-students are helped by the price going down.

And why is that going to happen? If you think about it-- OK, so, this thing has a-- this was the monopoly price here. So if I look at d(pi1)/d(P), that's what it looks like. d(pi)/d(P), if profit is concave, then d-pi dp is downward-sloping. It crosses 0 at p1*. And in this example, if I graph d(pi2)/d(P), again, it's downward-sloping. It crosses 0 at p2*. So then, if you ask, well, what if I just add this one and this one, I have one function that's going negative here. I have another function that's going negative here.

The optimal price, p*, is just the amount by which this one is negative equals the amount by this one is positive, so that the sum of the two first order conditions is 0. And as I've drawn it, here, this one is positive. And this one's 0. Here, this one's positive, and this one's negative. So obviously, the p* is going to be in between. And what I needed to do to say that is that, if these things are concave, then those first order conditions are both downward-sloping. And so, when they're downward-sloping, this happens.

Just one comment on elasticities, still, we just have this new demand, x of p, which is equal to x1 of p plus x2 of p. Your inverse elasticity pricing based on this aggregate demand-- and just to note, when you look at an aggregate demand, the elasticity in a market that's a component of two sub-markets is just a weighted average of the elasticities in the individual markets. And it's a weighted average of the elasticities in the individual markets, but where you're waiting by the share of the market.

So if there's a market that has more customers buying at that price, it gets a larger weight than one that's smaller. So it is the case that this is, again, another reason to think about why it's intermediate is, you have the high elasticity market. You have the low elasticity market. The combined market has an intermediate elasticity, which is a weighted average of the other two. So, therefore, it's in between, OK?

Well, this is what we think of as the normal case. I was assuming that pi* is concave. Profit functions are never concave, pretty much. What do profit functions look like? We normally draw them like this. But we ignore the fact that profits don't go to negative infinity as p goes to infinity. As p goes to infinity, people stop buying your product.

So what profit functions look like in practice is, they go like this. And then they asymptote to 0 as p goes to infinity. So they're concave part of the time. And then they become convex, as p gets large. And if you think about what happens when this is what profits look like, imagine what I have is the profit in the low demand market looks like this. And the profit in the high demand market looks like this, OK?

What do I get when I add those two together? I get something that goes like this. In some sense, after some point, when this one becomes essentially 0, this curve just coincides with this curve. So if I'm adding together one profit that looks like this and one profit that looks like this, there may be this bump over here on the left side, or maybe there isn't even a bump on the left side.

But the maximizer here can just coincide. If this one, imagine, no student is willing to pay more than $100 for a ticket to the A.R.T. or something like that. Then this just hits 0. When this hits 0, these things coincide. And what you do is you just abandon market 1 and monopoly price in market 2. So if you have that case where the concavity is not applying in the relevant region, you don't necessarily get the intermediate pricing result. You can get the just price to market 2 result.

If you just price to market 2, then, clearly, if when you price discriminate, you just price to market 2 and you eliminate the student discount, then, clearly, everybody is worse off. The monopolist makes less money because it's now no longer selling to any students. The students are worse off because they're not buying any tickets. They get zero surplus instead of getting positive surplus. So that's a case where there's a clear profit and social welfare and consumer surplus advantage to allowing price discrimination.

In general, the monopoly result is fully general. Monopolies always benefit from price discriminating because they could always choose not to. The monopolist could be told, OK, this customer is in France. This customer is in the Netherlands. They could just ignore that and set the same price. So the fact that they choose to set different prices means they're always better off discriminating. But whether consumer surplus and social welfare change, that's a different story. This normal case, we have the sense that some consumers would like you to stop discriminating. Some don't want you to stop discriminating. But that's not always true.

OK, a little more about welfare. So to think about aggregate social welfare of price discrimination, you have to think about two different welfare effects. So, one welfare effect of price discrimination, or not allowing price discrimination, is deadweight loss. How many consumers are there who would be willing to pay what it cost, who would like to buy the good at cost, who don't get served?

And then, the new problem we get in price discrimination is misallocation of goods sold, that when you sell the goods at different prices, there are high value people who are not buying the good. And there are people with lower values who are getting the good. So, for instance, let's think of a hypothetical world where, somehow, Apple has two markets to which they sell iPhones. They sell iPhones to people who want to use them as phones.

And they also sell iPhones to people who want to use them as paper weights. And somehow, they have this ability to restrict you when you buy it. I'm in the paperweight market or I'm in the phone market. And if you only buy in the paperweight market, they ensure that the phone actually won't ever work. It's still a working phone. They just somehow manage to block you from ever getting service on it.

If you imagine that world, it would be a funny world where you have some people who would like to pay $700 for an iPhone and who don't get to have an iPhone because the price is $799. And then you have other people-- or maybe it's a kids toy. It's an iPhone kids toy that's actually a functioning iPhone. And they get to buy it for $50.

And then, so you have someone who has this value for $50 for this toy who's getting to use the toy, and someone someone who's buying the phone at $700 isn't getting it. There's this gains from trade where, if the person owning the toy phone could hand it to the person who wants to use the phone, there'd be this big gain in social surplus. But you're going to get that any time you have a senior citizen discount.

Someone who doesn't even really want to go to the play is there and paying because it's only $15. And it was cheaper than going to the movie they would have rather seen. And then, you have someone else who would have liked to go who would be willing to pay $60, but isn't there. And there would be this-- you could just make this swap of, I'll trade you my play ticket. And I'll give you a ticket to the movie in the theater next door. We're both better off. Anyway, so there's that misallocation of goods utility.

A proposition just highlighting the second effect is, if third degree discrimination doesn't increase total output relative to uniform pricing, then it reduces social welfare. And this is going to be because if it doesn't increase output, it's not helping with deadweight loss. And so, if it's not helping with deadweight loss, it creates misallocation. And therefore, people are worse off.

So here's a proof. So imagine, we start from a situation where we have uniform pricing. And then we switch to discriminatory pricing. And I'm going to write delta x1 and delta x2 for how does demand change. So I'm going to write delta xi to be xi under discriminatory pricing, minus-- under discriminatory pricing, minus the xi under uniform pricing.

So here's a picture here. I guess I'll use this one instead of drawing a new one on the board. So if we had uniform pricing, this was going to be the price, p*. This would be the number of units sold, x*. And then, so this is the student market. These are the lower value consumers. When you start discriminating, you give them a price cut. You cut their price from p to p*. We get a positive delta xi. They buy this many units instead of this many units.

So a comment is, if I think about, what is the welfare gain in this market, the welfare gain in this market is, I've reduced deadweight loss by this much, this trapezoid here, because all of these units that were not formally getting purchased are now getting purchased at the lower price. And then this is the amount of surplus that's being generated on each of those units that's being purchased. So what I get is, if you think about this trapezoid, it's bounded by the big rectangle and the little rectangle.

And so that tells me that p* minus c times delta xi, that's the big rectangle, is bigger than delta Wi, is bigger than pi* minus c, times delta xi. So pi* minus c is the area of this little rectangle. That doesn't quite line up, but underneath. And then p* minus c is the height of the big rectangle. So the change in welfare is bounded between those two.

Notice, the same is true-- if there's another market where I reduce, I raise the price and I reduce the demand, again, the loss is going to be this trapezoid over here. And the loss with that trapezoid is going to be bounded between two rectangles, one being the p* minus c, which is now negative, the other being the pi* minus c, which is here. And with the negative numbers, the same inequality is still true. It's now, pi* minus c is the big negative number. This is the small negative number. It is in between the big negative number and the small negative number.

So if I just now add that across markets, summing over i, I sum this over markets, i. I get p* minus c times the sum of delta xi is bigger than the change in welfare, which is bigger than the sum over i of pi* minus c delta xi. And the left one is the one that's really involved in proving the proposition. If the sum of the delta xi*, if the sum of the delta xi's is 0 or negative, then this is 0 or a negative number. And it's bigger than the change in welfare. So the change in welfare is negative if output hasn't gone up.

OK, so some comments on here. This is a result about welfare. It's highlighting that there's that misallocation of goods sold. Obviously, output can go up. If output goes up, and output goes up by enough, then welfare can be higher with price discrimination. One example of that would be in the model where you only serve one market when you're discriminating. Then, obviously, output has gone up by the amount at which the students are buying things.

Equity concerns can go either way. We think about-- price discrimination, also, I said, is a transfer between groups, between the high demand group to the low demand group. What we think about that depends on the examples. So for instance, one case where we might think equity says we should discriminate is financial aid for undergraduate colleges in the US.

Colleges charge very, very different prices to different students. They charge much less to low income students than they charge to high income students. We tend to think of banning price discrimination there would be a bad thing because it would be raising the prices paid by all the low income students, who perhaps have a higher marginal utility of income.

Another example, though, that goes the opposite direction is health insurance. If you allow health insurance companies to discriminate, the people who should pay the high price are the sick people. And so, if you ban discrimination-- if you allow discrimination for health insurance, there, allowing discrimination will be charging high prices to sick people. Maybe we don't want to allow you to charge high prices to sick people for equity reasons.

OK, here is a corollary of the previous theorem, which I think is important to know. Suppose that demands in all these sub-markets are linear and that all markets are served under uniform pricing. Then third degree price discrimination always reduces welfare. So, again, why demand should be linear, I don't know. But it's a common benchmark that you might use as your first thought if you're studying a market. And this theorem says that, if the demands are linear, then third degree price discrimination reduces social welfare.

What's the argument for that? The argument for this is going to be that the sum-- in the linear model, the sum of the delta xi's is always equal to 0. And because the output is unchanged, then discrimination hurts social welfare. How do I show that? This is algebra, so I won't go through it. But this was something I think I've shown you over and over again when I've been drawing pictures on the board.

Suppose you have a linear demand curve. So this is my x of p. And this is p on this axis. I draw my cost in here. And then I ask, what's the monopoly price when I set-- what's the monopoly price when I have cost of c and I have a linear demand curve? It's exactly in the middle. So the rectangle that maximizes this is to set p halfway in between the price that's competitive.

Oops, what am I doing wrong here? Sorry, that's not what I wanted. I wanted x here and p here. So if I set price equal to cost, then this is the x in a competitive market. This is the 0. x monopoly equals 1/2 [INAUDIBLE] competitive. And then this is the monopoly price here. So if I put this on the x axis, I'm trying to figure out the monopoly price. It's always halfway in between 0 and what a competitive firm would charge-- would produce, OK?

So xi* is just one half of the competitive output. So if you sum across markets, what you just find is that the sum of the xi's is just one half of the sum of what you'd get under perfect competition in each market together. And if you think about perfect competition with all the markets combined, you set p equal to c, that's still one half of what you get with perfect competition in all markets.

So if I combine the markets together, I get one half of the competitive market as the total output. And if I separate the markets, I get one half of the competitive output in market 1, plus one half of the competitive output in market 2, plus one half of the competitive output in market n. That's just-- because the competitive output is always the demand at price equal to c, those two things are equal.

So the sum of one half of the competitive outputs, market by market, is one half of the competitive output in all the markets together. So with linear demand curves, because of this feature that you sell exactly half the competitive output, output is unchanged under-- output is unchanged if you allow discrimination. Therefore, discrimination reduces social welfare, OK?

I covered this theorem because I think it highlights a very general concern, which is, often, people, when they're doing structural estimation of models, they assume some theoretical structure, whether on the demand curves or on the utility functions. And then they estimate some number of parameters. And then they comment on something of interest. So this gives you an example.

Suppose someone comes to you and says, OK, I'm going to write a paper. I'm concerned about price discrimination in health insurance markets. I'm going to imagine that, instead of being allowed to discriminate and charge different health-related prices, I'm going to charge equal prices in all market. And I'm going to see what's the effect on social welfare of forcing the firm to charge a uniform price across all classes of consumers instead of discriminating by consumers.

And the student tells me, oh, wait, for simplicity, I'm going to assume that demand curves are linear in every market. Well, then, I already know what the results of the paper are. I know that the paper is going to say that charging uniform price has increased social welfare. And it's increased social welfare because it's implicit in this demand specification that you increase social welfare when you ban discrimination.

And I think it's just a very important thing to know, that if you're thinking about writing about a model to investigate some question, you have to know whether the answer is already determined by this functional form you've chosen. And again, why is linear demand curves not right? In part, social welfare is about how much utility people are getting beyond-- what's the consumer surplus that people are getting?

It gets very different if demand looks like this, and if there are some people who really like the product, or if it goes like that and there's nobody who likes it very much. In some sense, prices are determined by elasticities, which are just-- determine these slopes here. If you actually want to understand whether it's increasing or decreasing welfare, you'd have to be able to think about, can I estimate these things?

How much consumer surplus is here, relative to what the monopolist does, which is based on the elasticity at this one point the monopolist is choosing. And you need to be aware that this assumption is picking the answer for you. I don't know that it's not-- it may be a useful rule of thumb just to keep in mind. If, in the linear model, there's no effect on aggregate output, maybe we would think that aggregate output effects typically are small. But again, it's all based in the linearity there.

OK, what is general? So there's a recent paper by Bergemann, Brooks, and Morris that asked this question, what can we say generally about when you allow a monopolist to discriminate? Will welfare-- we know that, you allow monopolist to discriminate, profits go up. Will social welfare go up? Will consumer surplus go up or down? What are the trade-offs? And what Bergemann, Brooks, and Morris note is that it depends on what information the monopolist has about the consumers.

So let's think about the standard example where the consumer's values are uniform, 0, 1, c equals 0. If the monopolist has no information about the consumers, then it doesn't matter if you allow him to discriminate. He has no information on which to discriminate. What's the monopolist going to do? This is your standard monopoly pricing, p equals 1/2. Profit is 1/4. Consumer surplus is 1/8.

In their figure, that's this point here. Consumer surplus is 1/8. Producer surplus is 1/4, OK? One extreme is, what if the monopolist can perform first degree price discrimination? The monopolist knows everything. If the monopolist knows everything, then the monopolist charges each consumer their full value. We know there that the monopolist's profit has gone up tremendously. And we know that the consumer surplus has been driven to 0.

So here, if the monopolist knows everything about the consumers, then the answer to what happens if you allow discrimination is, profit goes up-- profit goes up. Consumer surplus goes down. Social welfare goes up because we're above the 45 degree line. But they note, in the paper, thinking the same kind of reasoning as Roesler and Szentes, there are other situations where consumer surplus and profit will also go up if you allow the monopolist to discriminate, that they both go up.

Here's their interesting and clever example. So suppose what we have is, we have this demand curve that comes from there being consumers with every different-- these sub-populations of consumers with every different value. Suppose what the monopolist gets is this signal that says, one of two things is true. Either the consumer is someone whose value-- so this is going to be my f of v, given my signal.

The consumer either has a value between 0.01 and 0.02, or the consumer's value is, with small probability, drawn from a uniform distribution on 0.02 up to 1. So this signal means there's no way their value is between 0 and 0.01. Very likely, it's between-- all consumers with values in this range will have that signal. And then there's some small possibility of a false positive. And it's actually a consumer with a value drawn from here, OK?

What does your profit look like? Well, your profit is increasing in price until you get to 0.01. As you drop your price from 0.01 to 0.02, per consumer profit, you'd lose many customers. But then, once you raise your price, you then are just making more money. Now, all these consumers have gone away. And you're making more money off these consumers as you raise your price. And you have something that looks like this.

In this sub-population, it's like your standard monopoly pricing. And what they note is that, by very clever choice of how many people to put in this band, I can make it so that this profit level and this profit level coincide. So I'm exactly indifferent between two prices, 0.01, or my standard 0.5 price. And so, if you give the monopolist a signal that says, there's a few consumers with this-- here's this signal. You get this signal, or you get this signal occasionally. And then, the rest of the time, you just get the regular signal. It's a regular customer.

Then, when you get this signal-- and then, you have the other signal that tells you-- the complementary signal would tell me that the consumer might have a value here, no chance to have a value here. And then they may have a value like that. So that's my complementary signal. It's either this signal or this signal, OK? Well, when I get this signal, my optimal price is p equals 0.5.

When I get this signal, I'm indifferent between 0.01 and 0.5. So if I charge 0.01 when I get this signal, I've just created consumer surplus. And I haven't affected profits at all. So what I'm able to do is move along this line and just keep profits exactly at the same thing I get charging 0.5 to every one, and give consumer surplus. And then you can think about this. I could do the same argument, adding more and more such signals. I can do a 0.02 to 0.03 signal, a 0.03 to 0.04, 0.04 to 0.05. I can give the monopolist many, many, many signals, and just, by doing that, get the monopolist to have more and more groups to which it offers a discount and increase consumer surplus up to some point here.

It turns out, I can also-- so that's their clever construction for how I get from here, going in this direction. It turns out I can also do the same kind of clever construction and go in that direction. I can make the monopolist indifferent between 0.5 and a very high price. And then, once the monopolist is indifferent to 0.5 and a very high price, they start charging the high price to all these. There's a special signal that tells you that.

There's some chance, there's this excess probability that the value is up here. And otherwise, it's uniform. When you get those signals, they're in between charging the very high price in selling to these consumers only, or setting the 0.5 price. And so, by using other signals, I can keep the profit-- here, the profit went up as consumer surplus went to 0. Here, I can have the monopolist profit stay the same and drive consumer surplus to 0.

And their general result is, by this choice of clever distributions of signals, what they show is that, really, what price discrimination does to consumer surplus and to social welfare is just totally indeterminate. If the monopolist has some types of information structures, it always raises the monopoly profit. But it can raise or lower consumer surplus. And because it can go above or below this 45 degree line, it can raise or lower total social welfare. And it's all a function of, what information does the monopolist have? And is it the type of information that makes the monopolist want to charge low prices, or is it the type of information that wants the monopolist to charge high prices?

AUDIENCE: So this result goes away if the monopolist is allowed to design their own information scheme?

GLENN ELLISON: That's right. So this is very much thinking about the-- the monopolist is given a scheme that, it's given a scheme like you either know that the person is astute, because it's third degree discrimination. You need to be able to price based on the signals that you're given. You've got demand in Mississippi, demand in Kansas, demand in Boston. That's all you know. Or you've got student demand. You've got non-student demand. So, obviously, the monopolist wants perfect information and wants to go here. And so, the more useful information the monopolist has, the higher it's going to get up here. But if the sub-populations are what they are, it can go either way, yeah.

Yeah, because I think, on the monopolist's side, the real intuition is, you want to know everything and do first degree. And if you can do that, that's what you're going to do. But as you're getting there, it's possible you're going there this way or this way and reducing welfare. It could be, you're going there this way and increasing welfare on the path there.

OK, so the final thing I'm going to do today, well, actually, next to last thing, is second degree discrimination. So now, we're back in the case of the iPhone 128 and the iPhone 512, where the monopolist can't observe the consumer preferences and is relying on the self-selection to get the high value consumers to buy the high-end good. So I'm going to have consumers have some unknown type, theta, that's unobserved by the monopolist. They get utility, v of q theta if they buy a quality, q, good, and their type is theta.

And they pay a total price of T. I'm going to assume that the utility is that everyone agrees that quality is good. The higher theta consumers have a higher willingness to pay for any good. Cross-partial of v with respect to q and theta is positive. So the higher theta, people have a higher valuation and a higher marginal valuation for quality. And I'm going to assume that there are diminishing returns to quality. So second derivative of v with respect to q is negative.

I'm going to, and this is really a labeling of quality, I'm going to assume that the cost of producing a quality, q, good is c times q. So you can think of that as just determining how we call-- we don't want to call a 128 versus 512 four times as good. The q is defined in terms of how much the cost of producing that quality is. For now, I'm just going to assume there are two types. And I'm going to write vi of q for the-- so I'm going to write v1 of q and v2 of q for what the utility function is for the low types and the high types.

Again, if the thetas were observable, first. this would be first degree discrimination model. I would just, for every consumer, sell them the efficient quality for them and charge them their full surplus that they get from it. With theta unobserved, this isn't going to work. In the perfect price discrimination scheme, the type theta equals 2 consumers get zero utility when they buy q2* at a price of T2*. These were the first degree levels.

So they get zero utility from buying the first degree packaging for them. But they're going to get positive utility from buying q1* at a price of T1* because the low types get non-negative utility from buying a quality q1* good at price T1*. The theta 2 types always get more utility than the theta 1 get from any utility. Therefore, the theta 2 types would be better off not buying the package designed for them, and instead-- so in some sense, if the business class tickets were the full willingness to pay of the businessmen, they would go to fly in coach and get a positive surplus from sitting in a coach seat at the coach price.

OK, so what do you want to do with theta unobservable? This is your standard screening problem. You have two types. You offer a menu that has two options. And you let the consumers choose between those two options. So consumers are given the option, you can either buy a quality q1 good at a price of T1 or a quality q2 good at a price of T2. Monopolist maximization problem, I'm mazimizing. So I have a equal masses of consumers of each type here.

I'm getting the price I get from the low types, plus the price I get from the high types, minus the cost of producing all the goods. And then I have these four standard constraints like you're used to seeing in mechanism design problems. The IR constraints, that stands for individual rationality. The constraints say that the low type consumers have to be getting at least as much utility from the good they're buying as what they're paying for it, as do the high types.

And then, the IC constraints, the incentive compatibility, the incentive compatibility constraints say that the utility that the low types get from buying the package designed for them is at least as good as the utility they would get if they bought the package that was instead designed for the high types and paid the high-type price. And similarly, the businessmen have to be getting at least as much utility from sitting in business class as they would get if they instead sat in coach, got the coach quality, and paid the coach price.

So each group has to be willing to buy the package that you're selling to them in equilibrium, relative to the other group's package. The first step in solving problems like this is always to figure out which constraints are binding and which constraints are not binding. Here, one observation is that the IR2 constraint is not going to be binding. And the argument here is simply that the students are willing to fly on the plane in coach. The IC constraint says that the businessmen have to prefer what they're getting to flying in coach. They have positive utility from that.

So therefore, they're automatically willing to buy because they would get more utility from the coach class seat. And then, I would say the second thing is just a guess, although I'll show you, it's true. IC1 seems unlikely to bind. If you're thinking about, what's the airline's pricing problem, you never have the problem of, the students are going to buy the $3,000 ticket if you're not careful. The problem is keeping the wealthy people from buying the cheap product, not keeping the poor people from buying the expensive product.

So once I eliminate those two, the simplified problem is, I'm maximizing my profits subject to just two constraints. The low types have to be willing to buy the product I'm selling them. And the high types have to prefer the expensive product I'm trying to sell them to the cheaper product I'm selling to the low types. So if you look at this problem now, so I'm doing this constrained optimization, I want to make the Ts as big as I can. What are the constraints on making the Ts as big as I can? So for T1, if I raise it, that's good. It makes this constraint easier to satisfy. So the only reason not to raise T1 is because this constraint is going to be violated. So clearly, in equilibrium, I'm going to set T1 equal to v1 one q1 and have this constraint be binding.

And then, also, I'm going to want to raise the price to the high types as much as I can. What prevents me from doing that is only that T2 appears in this constraint. And so, I'm going to make T2 as large as I can before this constraint is violated. And so, this is going to also hold with equality. And that means, if I bring it over to this side, T2 equals T1, plus v2 of q2, minus v2 of q1.

So in some sense, this is just like in the durable goods problem, where people could buy the good in the second period. So I can charge them for getting the good in the first period instead. Is there incremental utility for getting it one period earlier? Again, what I can get from selling them the higher quality good is what they would pay for the low quality good, plus the incremental utility that they get from the quality upgrade. That's the most they're going to be able to pay.

And so, that's v1 q1, plus v2 of q2 minus v2 of q1. And so, once I've made these two substitutions, this for T1 and this for T2, I now just have an unconstrained maximization problem. What's the answer to that? So I've just-- assuming the v's are concave functions, I just take the first order conditions for this to maximize. First, let me differentiate with respect to q2, which only appears here and here. I differentiate with respect to q2. And I get dv, d(v2)/d(q) equals c.

So that just tells me that the q2 with second degree discrimination is equal to the q2* that you do with first degree discrimination. Or with the monopoly Apple equality theorem, the high types get the optimal product. And I think the sort of thing-- sense you get. Like, sometimes, I know people have this sense that there's excess consumption, or excess quantity.

Like, I walk through that first class cabin on the airplane. It isn't that obscene that these people are drinking drinks from glass glasses and using metal silverware, or whatever. And you just have to think to yourself, wait, no, actually, I'm relatively low in the income distribution. And I have metal silverware at my house. And I drink out of glass cups instead of paper cups. It's just, that is, given the cost of providing metal, it seems like that's a totally reasonable price to be paying for people, or whatever. So you give them the optimal quality product and charge them for it.

And then, what about q1? I differentiate this with respect to q1. I get a d(v)/d(q1). Again, it's equal to c. But then, I'm going to have two other terms, derivatives, d(v1)/d(q1) and d(v2)/d(q1). And so, what you get is d(v1)/d(q1), moving everything to the other side, equals c, plus partial of v2 with respect to q, minus partial of v1 with respect to q.

Remember, the higher types have a higher incremental willingness to pay for quality at the same quality level. So this is greater than c. So the distortion is, I distort quality down for the low types and provide them with a suboptimal quality. If you've done that calculation on the airplane, I can't believe I'm sitting here in coach. And I can't open my laptop whatsoever. And my knees are pressed in the seat in front of me. Couldn't we just have moved the seats one inch further apart, take one seat out of the plane? I would be willing to pay this much more. They would make more money. Why don't they do that?

Well, the answer is, yes, they could make more money in the coach cabin doing that and making the seats more comfortable. But then, the businessmen would be willing to sit in coach. So you have to make those coach seats uncomfortable enough that the businessmen will not sit in them. Your seats are the ones that are distorted. Theirs are where your knees don't bang into the seat in front of you. That's optimal quality. And so, what we do is, we distort quality down for the low types.

And so, this is a new inefficiency of second degree price discrimination. We have all the old efficiencies, inefficiencies. Our new inefficiency, again, is suboptimal product quality. And the sub-optimal product quality is done-- it's because they don't have the ID to say, I'm rich or I'm poor, which seat I'm going to sit in. We have to have this extra distortion to allow them to screen without that signal, OK?

What's the other comments? So the high types are getting positive surplus. The low types are not. We sometimes call that an information rent. Welfare is worse than first degree discrimination. If we could give the students their optimal seat pitch and charge them the optimal price, they would be better off. But the comparison with what would happen, for instance, if we require you to have only a single good and not have two different coach classes, that's an ambiguous distinction.

So again, the fact that you may be able to create these low quality seats and serve people who would not be in the market if there were no low quality seats, social welfare can be better than if you had the only first class plane, for instance, which is what you might get if you didn't allow discrimination. OK, I would say, modern IO theory papers tend to work with a continuum of types model instead of a discrete-type model.

These days, I have this awkward thing that Alex, when he teaches [14.]124, does a really good job on the continuum-type screening models. So I'm not going to cover continuum-type screening models. It's in the slides that I have here. [INAUDIBLE] is going to cover it in recitation on Friday. I think, I haven't looked at it this morning, but I think there's also a 2020 me discussing this material on Zoom video if you can't make recitation on Friday, and see me discussing the continuum of types price discrimination in a COVID world. I think that's on the website, too.

But I mean, the general intuitions are just that, if you have three or four or five types instead of-- or a continuum of types instead of two types, you need to stop these guys from copying these guys buying the product. And then, for these guys, you need to stop these people from buying this product. You need to stop these people from buying this product.

And so, in some sense, you get these cascading distortions, where any time I give something to the lowest type, the next lowest type gets more utility because they have to be given utility so they won't copy the lowest type. And then, because they're getting a better product, the next type up have to get a better product, and so on. So you get these places where you have very severe distortions, which is that the further you go down in the distribution, the more reasons you have to distort the package being given by the low, low types, because you're giving money to every higher type when you give them something.

And then, a second intuition is, if there are types that there aren't a lot of, you want to give them a very bad deal because anything you give them, the fact that it's available raises the utility of the higher types. If there aren't very many people in that group, then you really want to discriminate. You want to give them a terrible product because the profit you get from giving them optimal quality is tiny when there aren't many people there. So people who are in a small fraction of the population, they also get really bad service on quality.

But anyway, so there's a video on the website. I think what I want to do to jump ahead at the end is, let me just talk about bundling. So let's think about, you have a firm that's selling multiple products. What's the best way to sell multiple goods instead of a single good? So, for instance, Apple doesn't just sell phones. They sell phones, and chargers, and earbuds, and apps for the phones.

Netflix doesn't just sell movie by movie. Netflix sells whole big bundles of movies. So, observation from this literature is that, often, if a firm sells multiple products that are used together, they're better off doing some kind of bundled pricing. Why do you do some kind of bundled pricing? I put it in this lecture because often, the reasons why you want to bundle goods together have something to do with price discrimination.

So, the first example, due to Stigler, is, imagine you have consumers who have values for the two goods that you sell which are negatively correlated. So the example that Stigler gave is, suppose you have these types, theta, that are uniformly distributed on 0, 1. And these are about their preferences over movies. And so, a type theta consumer gets utility theta from watching some kind of superhero action movie and gets utility 1 minus theta from watching a romantic comedy.

And people's preference over these things are just exactly negatively correlated. So if I'm a far right superhero movie person, I'm also the far left romantic comedy utility. And if I'm the far right romantic comedy utility, I'm the far left superhero quality. If Netflix had to price these things separately, and we have Netflix action/adventure movies and Netflix romantic comedy movies, what's going to happen? Well, if these are the values, it's going to set p* equals 1/2 for each service.

Every consumer is going to buy exactly one of the two services. And so, Netflix ends up earning profit 1/4 on each. I set a price of 1/2. I sell the 1/2 to the consumers. And that's the profits that I get. Suppose they instead follow this pure bundling thing, where we'll have both the action/adventure movies and the romantic comedies on the same service. Well, then, what everybody can do is, they can both watch their superhero movies. And they can watch romantic comedies, because they're both available free on the service.

What utility do they get from doing that? They get utility theta plus 1 minus theta equals 1. And so, because they get utility 1, you can just set the price of the bundle equals 1. It's fully socially-efficient. We extract all of the social surplus. But one way to think about this bundle of what we're doing is, we just have the people-- we know the people who are buying the superhero movie channel have low values for the romantic movies. So we just want to charge them a lower price for that. And this combination of charging them a high price for one and a low price for one ends up being something you can do perfectly by bundling products together. So that would be one argument for bundling.

There's a very nice paper by McAfee, MacMillan, and Winston that notes that Stigler gave this example making you think negative correlation is important. Actually, the same thing works with independent valuations. So suppose your values for romantic comedies and for superhero movies are completely independent draws. Then it's still the case that the firm would be better off bundling, not pure bundling like in Stigler, but with mixed bundling. You offer three packages.

You offer the superhero movie channel, the romantic comedy channel, or Netflix total. And this mixed bundling scheme of offering the bundle and the two products separately strictly dominates offering the product separately, when you have independent valuations, OK? In some sense, it's an obvious statement that the mixed bundling strategy, this offer, p1, p2, and p-bundle, is at least weakly better, because you could always offer good 1 at p1*, good 2 at p2*, and price the bundle at p1* plus p2*, the bundle, in some sense, superfluous. But it doesn't hurt the monopolist's profits.

But what McAfee, MacMillan, and Winston show is that, in fact, it strictly raises profits. And so, I'll just go through this as quickly as I can. So, start from the situation where you have p1*, p2*, and the price of the bundle. Then what's going to happen is, these consumers for which v1 is large are going to buy good 1. These for whom v2 is large are going to buy good 2. And then people here are going to buy the bundle.

Imagine that what the monopolist does now is, it raises the price of good 2 by epsilon and leaves the bundle price unchanged. Who is going to react when you raise the price of good 2, when you leave the bundle price unchanged? Well, these consumers who would have formerly bought just good 2, now that good 2 has gone up in price, they're going to switch from buying good 2 to buying the bundle.

And so we get this first order gain in profits from all of these people who were buying just good 2, now also buying good 1 from me, which has a positive markup. What do I lose? I lose these people who would have bought good 2 at p2*, but now don't want to pay p2* plus epsilon because it's not worth it to them. But we know that d(pi2)/d(p2) evaluated at p2* equals 0.

So the fact that I lose these customers and charge a price here, I was the first order-indifferent between this price and this price anyway in this market. So I've got a second order loss from raising the price, a first order benefit from cutting the price-- sorry, from raising the price and getting them to buy the bundle instead. And therefore, I'm strictly better off by bundling.

And so, that's the basic argument in McAfee, MacMillan, and Winston, is that mixed bundling is obviously weakly optimal. And if you just think about starting from that weakly optimal thing, and then just raising the price of the two goods individually, that's always better for you. So the fact that you can bundle and have people have this option to switch to the bundle, which lets you sell extra goods is good for you. And so, you do generally want to bundle if you're in this case, where values for the two goods are independent. OK, so I am out of time. So why don't I just stop there?