[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: Welcome to class. So this is 14.271. I should say that syllabus has a bunch of basic information. You'll notice the course is listed as co-taught by me and Tobias. COVID permitting, I'm going to be doing most of the lecturing in this class. I think the plan is I'll be doing 21 of the lectures, Tobias will be doing four, Jean Tirole will be doing two.

And it's a basic course in industrial organization. I think these days, my primary responsibility in this class is to teach students who are going to write their thesis in industrial organization what they need to do to write an IO thesis. But the course has also traditionally been useful for many students in other fields.

I think that's for two reasons. One is IO has always been at the forefront in economics and linking theory and empirics together. And then also in recent years, IO has developed a lot of techniques for what's referred to as structural, empirical industrial organization, which have become very, very popular in other fields. Anyway. So my goal here is really to teach the basics of IO theory and how IO people think about empirical work, which is sometimes different from how empirical work is done in different fields.

The way I'm going to do that is you'll see on the syllabus we've got 12 topics or whatever we're going to cover. Almost all of them are going to get covered in two lectures. And typically in the first part of the semester, the pattern is going to be on Wednesday, I'm going to explain the theory of a topic. I'm going to go over the basic theory on a topic. And then on the Monday, I'm going to go over the empirical work on a topic.

And my basic view of IO is IO is just an inherently theoretical field. There are many fields in which empirical work and natural experiments testing policies are what you want to do in studying that field. In IO, the questions we get asked are should Facebook be required to divest itself of Instagram? Or is it OK for Google to be paying Apple $12 billion a year to make Google the default search engine on iPhones?

And these are not things that can be addressed with natural experiments. You can't create 100 different worlds and break up Facebook and Instagram in half of them and not break up Facebook and Instagram in the other half of them and then do a comparison of means 10 years later and see which was the right decision. It's a one time decision that will get made or not. And so you basically make decisions like that. You need to have a theoretical understanding of how the world works and how firms competing with other firms work and what do you think would happen on the two different market structures. And then you may never find out if you were right or wrong.

While I think IO is inherently a theoretical field, almost all students getting good jobs in IO in the last 15, 20 years are doing empirical IO. And in part, this is because what happens is people in IO make up models. And the question is you can make up models that predict all kinds of things. Then you have to ask which model's assumptions are correct in terms of how people behave, how firms behave, and how large are those effects empirically.

And so what typically happens, there's a lot of interplay in IO between you develop a theory, you then do empirical work that tries to examine the theory in places where it should or shouldn't work, and you find out how important are the effects the different models predict. Are the assumptions and models good or should we be looking at what-- in some sense, in any good IO empirical work, something doesn't fit. And then what doesn't fit gives you the idea of how do you refine the model, what else should be added to the model to get a better predictive model.

So let me go ahead and get started. And I'm going to get started-- I should say normally when you design a class, you want to have the first class be really exciting and have something intriguing for students to get hooked on the class. That's not the case with IO. I have to start at square one. I have to start with monopoly pricing. I'm sure many of you have seen monopoly pricing in your undergraduate intermediate micro or principles classes. But I'm going to start with the beginning with monopolies pricing and then build from there.

So anyway, first class. So this is 14.01. Monopoly seller of a single product. In this class, I'm generally going to be putting quantity on the x-axis and price on the y-axis. So a monopolist has this downward facing, this inverse demand curve p of x if it sells x units. Monopolist has a total cost function c of x. So if they produce x units, the total cost of production is c of x.

What does the monopolist do? You've seen this many times, you maximize over x, x times p of x minus c of x. I take the first order condition and I get this familiar equation. Marginal revenue minus marginal cost equals 0. So the monopolist sets marginal revenue equal to marginal cost.

Why am I starting here? Well, even in IO, we do things a little differently. The way people typically think about this first order condition is instead of putting marginal revenue on the left and marginal cost on the right, I'm going to take price minus marginal cost and move the minus x p prime at x to the other side. And then I divide both sides by p of x and I get price minus marginal cost over price equals this.

And what's over here on the right side? p prime at x, x over p of x. So you can think of that as x over p of x, 1 over p prime at x. And then if you think about the inverse function theorem, 1 over p prime at x is kind of like dq / dp. Because p is the-- if you think about p as the inverse function of q, so this is like x over p prime at x. And so if I divide this, dq / dp times p over q, this is the elasticity of demand.

So what a monopolist does is it sets its price so that price minus marginal cost over price is the negative of 1 over the elasticity of demand. And I start here because in macroeconomics, they seem to have this thing that when demand goes up, price goes up, and when demand goes down, price goes down. That doesn't happen in IO. When you think about firms with market power, pricing is based on demand elasticities, not demand levels.

And on this left side, this is something that in IO we call it the Lerner index. Ignoring the marginal versus average, this is like what accountants would call the gross margin. This is the markup that you're charging as a fraction of the price. So if you charge double your marginal cost, that's a 50% markup, because 50% of your price is markup, 50% is cost. So this says the Lerner index is negative of 1 over the elasticity of demand.

So if, for instance, the elasticity of demand is-- if epsilon is negative 2, then this would say you do price minus marginal cost over price is 1/2. So if you're a monopolist, your elasticity of demand is 2. Your margin is 50%. You double your marginal cost when you think of your monopoly price. If your elasticity is 4, you do your 25% Lerner index. Lerner index is 1/4, which means you're raising your price by 1/3 over cost.

In 14.01, you certainly also did a diagram like this. So what happens? Firm sets marginal revenue equal to marginal cost. They have this monopoly price. They earn a profit on all the units they sell at a cost that's greater than the cost of producing them. And so this is the profit. This is the difference between the consumer willingness to pay for each unit and the price that they've paid. So this is consumer surplus. This is deadweight loss.

Typically, a lot of IO is about market power is a bad thing. We want to prevent market power. Why do we want to prevent market power? Classic explanation for why you want to prevent market power is this triangle. You have this deadweight loss is a social inefficiency. People are willing to pay more than the cost of producing the goods, and yet those goods don't get produced and sold to them. That's socially wasteful that we have these opportunities for gains from trade that aren't realized.

While people often focus on deadweight loss, there are a lot of other reasons why you might think monopolies are also a bad idea. So many of these are harder to actually explain economically, to realize I think they should be recurring. But one thing you hear people talk about is productive inefficiency. Certainly the question is, are monopolists going to produce things at the lowest cost possible? Your basic intuition should be yes. Monopolist gets all the extra money they make from cost savings. So monopolists ought to want to produce the most efficient technology possible.

There are lots of discussions, though. If you think about the old Soviet Union, they had factories making things out of titanium that could have been made out of steel and that cost 5 times as much as they cost in the West. Why did that happen? Well, one thing that you hear is that if you have a firm that's a monopoly, they don't have anyone to compare themselves to. They don't realize there are better production technologies. Employees can be lazy in figuring out how to produce things and no one ever catches them. So productive inefficiency could be a cost of monopoly.

Quality distortions. Again, this is more an informal complaint. But people look at Microsoft Windows and are often thinking of 30 different criticisms of Microsoft Windows and why does it suck in so many ways. Why can't Excel do this? Why can't Excel do this? Is it because Microsoft is a monopoly and has no incentive to improve the quality of its products because you've got to buy them? You've got to put Windows on your computer anyway, even if you don't like it.

Rent seeking is an important thing. Sometimes we have monopolies for regulatory reasons or for capacity constraints. So, for instance, there's a limited number of licenses to provide cell phone service in every city. In the old days, when cell phone service first started, these were actually-- when the United States first adopted cell phones, there were going to be two licenses in every city.

And they had what people refer to as a beauty contest to allocate those licenses. That is, firms didn't have to pay for them. You just had to apply to provide cell phone service and the government would pick the firm that seemed best, socially efficient to provide the cell phone service, and they would award those licenses to those winning parties. Often these winning parties were perhaps not the ones you would think.

So, for instance, I don't know if the story is true. There's a story that in Washington, DC, one of the two licenses went to Katharine Graham, who was the editor of the *Washington Post.* Why would she know more than anyone else about running a cell phone company? Well, she ran a newspaper that would write stories about the politicians. So when they're thinking of who to give the license to, they give it to her.

But in general, in many situations, you're going to find that if there are monopoly rents to be given out, someone has the power to give out the monopoly rents. And then when someone has the power to give out the monopoly rents, those people may take bribes from the people who are giving them. It may be that you have something akin to a political campaign where you're doing advertising or lobbying, and you have all these socially wasteful expenditures to try to be the one who earns these profits.

If you imagine a bidding game where if everybody is going to give a bribe to the politician and the biggest bribe wins, what's your bribe going to be? You're willing to pay a bribe equal to the discounted value of all future profits that you're going to earn. In that game, you'd expect all of the profits would also then get dissipated socially.

Well, actually, now if it's a bribe, you believe in efficient markets, it's just a transfer. And so it's just a transfer of wealth from consumers to the one person who's taking the bribe. But if the bribes are given in a socially wasteful way, all of this could also be a loss of monopoly power.

And then my final comment on bribes brings up this sort of equity concern. Another problem with monopoly is that in a competitive market, the consumers will be paying this. In a monopoly, they're paying this. Consumers paying more. It's a transfer from consumers to firms. If you believe that the consumers have a higher marginal utility of income than the shareholders of the firm, that's also a concern we should be worried about.

I should say that in IO, I mentioned that today, I'm never going to mention it again pretty much. Jean may talk about it when he gives his guest lectures. But in IO, all of these income effects you see in public finance, whatever, those are all going to disappear. So in IO, pretty much everybody has quasi linear preferences and everyone has preferences that look like v minus p, where everybody's got a marginal utility of income of 1 that's constant and everyone has some value from buying the good. And we just do additively separable utility and don't really worry about income effects and stuff.

And obviously that's not to say that equity concerns and transfers across consumers aren't important, but it's just I think people in IO have made the decision that you don't want to be every single model talking about the incentive to distort what you have the firms doing in order to benefit the consumers with lower marginal utilities of income instead of higher marginal utilities of income. So we just do this, and you can think about equity concerns on top of it. But pretty much all semester I'm going to have that kind of utility.

OK. That's your basic monopoly pricing model. What are some of the other things you need to know? First of all, most firms in the real world sell multiple goods instead of a single good. So the way I'm going to think about this is simplest example. Suppose the monopolist has two goods to sell.

Now I'm putting price as the argument. So demand for good 1 is a function of p1 and p2. Demand for good 2 is also a function of p1 and p2. Let me stick with constant marginal costs here. So the total cost of producing x1 units of good 1 and x2 units of good 2 is just c1 x1 plus c2 x2.

So what's the monopolist's problem? It's now maximizing over the two prices. Price 1 minus marginal cost 1 times the number of units of good 1 plus p2 minus c2 dx times x2. What do the first order conditions look like? Well, if I differentiate, think about differentiating with respect to the price of good 1, if I differentiate with respect to p1, I'm going to get p1 minus c1 d(x1)/d(p1). Plus this times the derivative of that, which is x1 of P1 P2. But now I'm going to get an extra term that I wasn't getting before, which is d(x2)/d(p1) times p2 minus c2.

So one way to think about what we're getting is if I now do the same thing, if I divide by x, I get price minus cost plus x1 over d(x1)/d(p1) plus p2 minus c2 times d(x2)/d(p1) over d(x1)/d(p1) equals 0. So this is the same first order condition I had before, but with the addition of this extra term.

People in IO refer to this as the diversion ratio. This is when I raise the price of good 1, I'm going to sell fewer units of good 2. When I sell fewer units of good 2, sorry, when I sell price of good 1, I'm going to sell fewer units of good 1. But when I sell fewer units of good 1, that's going to affect my sales of good 2.

And so it can either raise my sales of good 2 or lower my sales of good 2. But I have to take in that account for every unit of goods that I don't sell. It gets diverted possibly to a sale of another good. And I have to take the profits I would earn on those goods into account when I'm doing the first one.

And so this gives me a different markup formula. Price minus marginal cost over price is minus 1 over epsilon. The elasticity of the good plus this other effect, which comes from thinking about the effect of my change in price of good 1 on sales of good 2.

I should say before I go to the next page, one way you can also think about this is just you have to think about your pricing as a system. So, for instance, if you're-- the worst thing in the world you could do is imagine you're advising Apple on selling both blue iPhones and silver iPhones.

And you do a study of the price elasticity of blue iPhones by raising the price of blue iPhones by $20 and lowering it by $20 and seeing what that does to sales, holding all other prices fixed. You'll find that the blue iPhone is incredibly price sensitive price, because people switch and buy the silver one.

That does not mean that you want to have a really, really low markup on the blue iPhone. You have to think of it as a system and what's the equilibrium solution to blue iPhone, silver iPhone, red iPhone, green iPhone, whatever. And think of all those things as a system. If you look at each one individually, you'll get it terribly wrong.

And so in general, any time you're doing pricing of multiple goods, if there are very strong diversion ratios between them, you always want to make sure you're thinking of it as a system, not in terms of each one individually as its epsilon. But when these diversion ratios are smaller, you can think of it as it is 1 over epsilon pricing with some increase or decrease based on whether there is an effect of changing sales on the other products. OK, anyway.

OK, so what does this formula look like? So my price is inverse elasticity minus price minus cost over price d(xi)/d(pi). If you think about the normal case, in the normal case, goods are sold at positive markups. So this would be a positive number. This would be a positive number. The d(xi)/d(pi) is a negative number. So this is negative. There's a negative sign. This is positive. So you can think of the whole thing in blue as a positive number normally.

And so whether you price above or below the inverse elasticity is typically going to depend on whether this is a positive number or a negative number. Standard case, we're dealing with a monopoly. Pricing is often the substitutes case. The substitutes case is when I raise the price of the blue iPhone, sales of the silver iPhone go up. These goods are substitutes for consumers.

So when that's the case, margins for both products are positive. And your margin exceeds the single product Lerner index formula. And the amount by which it exceeds that single index Lerner formula depends on what is dx minus i over pdpi. How large is the diversion ratio? This divided by this is the diversion ratio. So the larger is the diversion ratio, the larger is the markup you're earning on the other second product you're selling, the more you exceed the Lerner index formula.

The other more interesting case is this complements case. So sometimes you have firms selling complements like hamburgers and French fries at McDonald's. If your price of hamburgers is lower, more people come to McDonald's, more people buy French fries. If your price for hamburgers is higher, there could be an income effect where people feel like I've spent too much already. I don't want to also buy the fries. But typically it's like the effect is going to be people don't come to your restaurant and then they don't buy fries at all.

In the Apple case, if I don't sell them this iPhone, I'm not going to sell them my earbuds that I sell for a huge profit. If I don't sell them my iPhone, I'm not going to sell them the 17th charger that they've had to buy from me because I keep changing the way my charger and connector works. And so the complements case is the higher is the price of my iPhone, I sell fewer. If I raise the price of the iPhone, I sell fewer earbuds.

And so in this case, now this becomes a negative number. You set price below the elasticity of demand. Very easily if this diversion ratio is large, you can often think of the diversion ratio is in some sense 1. For every iPhone I sell, I don't sell, that's exactly one pair of earbuds I don't sell or that's one charger I don't sell. So if this is a large number, if this diversion ratio is 1, the margin is high in the competing good, you can very easily end up with pricing one good below cost or even having a negative price.

You may have thought you go into the store to buy an Inkjet printer and it's remarkable, but I can get this phenomenally complicated piece of electronics with all kinds of chips in it and hardware and precision stuff, and it's only $99. And then you get really pissed off when the thing runs out of ink seven days later and it's $59 for ink.

And you're thinking yourself, how can it be that it only costs $99 to produce this printer and yet it costs $59 to get a bottle of ink this big? You would think the ink should only cost $0.50. And the answer is yes, obviously, yes. The ink only costs $0.50. But they're pricing the printer way below the cost of producing the printer to get you to buy the ink.

I guess newspapers, these examples no longer resonate. There was a time when people had these things called newspapers. And instead of reading news on your phone, you got this big piece of paper delivered to your house every day. The way newspapers used to work is the newspaper would only cost $0.50. And if you think about how much does it cost to produce and get a newspaper to your house every morning, the answer was way more than $0.50. It would cost $1 or $1.50 or something like that.

How did newspapers work? Newspapers worked because they sold them for $0.50, even though it cost $1.50 to produce it because they were full of ads. And they were full of classified ads that now go to wherever. They were full of real estate ads that are now on Zillow. They were full of ads for local car dealers. They were full of ads for all kinds of things. Whatever, eBay, everything that was sold on eBay was sold through the newspaper.

The newspapers were selling the newspaper below cost because the cost of printing an ad was essentially 0. And so the cost of producing an ad was 0. They'd charge you $100 or $1,000 to list your house in the newspaper. Price minus cost was huge on the ads that made the margin on the newspaper negative.

In some industries, it even becomes there's a 0 lower bound that becomes important. Google in some sense would like to pay you to use Google. Every time you do a-- Google's ad revenue is something like $0.06 or $0.08 a search. Every time you do a search on Google, Google is making money. The marginal cost of running a search for you on Google is very, very small.

So when Google does the price search, it costs them a thousandth of a cent to run the search for you. They're making $0.08 profit on every search. They want to set the price of having Google on your phone negative. That's why Google is willing to pay Apple billions and billions of dollars to put Google there and make Google a default. In some sense, they want to have a negative price to get as many people to use it as possible. Any questions on multiproduct pricing?

So product quality. So as I said, people often get frustrated with their Microsoft and other products and think this thing is so bad because they're a monopoly. I have no choice. Maybe they get frustrated with the Sloan cafeteria and are like, this place is closer than everywhere else and yet still people walk across to Brothers Market to get their lunch anyway, because it's not necessarily the price of the food at the Sloan cafeteria, although it's also high, but does Sloan cafeteria have the option to set the optimal product quality?

Basic reasoning is going to be, in some sense, while we may think the monopolists often provide low quality products, at least the profit maximizing theory tells us, they ought to do pretty well on this dimension. So let's suppose you have a monopolist. It has a product that's a variable quality s. And s is a choice variable as well as the price. So suppose there's a constant marginal cost c of s.

So you could think of this as either the quality of the food you're providing at a lunch or s could be the quality of the camera in an iPhone. The unit cost of producing this lunch or the unit cost of producing the iPhone is going to be c of s. And the higher quality you make it, the more expensive it's going to produce. And then I will do this as a constant marginal cost. So if you produce 100 iPhones, it's just 100 times the cost of producing one. If you produce 100 lunches, 100 times the cost of one lunch.

Something we're going to do all the time in IO is take consumer heterogeneity very seriously. So I'm going to think about massive consumers who have types theta distributed uniform on 0, 1. The uniform 0, 1 here is without loss of generality. I do that a lot. It's just I index consumers by their percentile within the distribution of consumers. Here it's indexing within their distribution of consumers in terms of how high their valuation is for the good that people are buying.

So consumers are going to get-- consumers are going to either buy one-- they either want one lunch or one phone. It's a unit demand case. And they're going to get utility v of s theta minus p if they buy one unit of a quality s good and their type is theta. And I'm going to assume that we've ordered the qualities and the types so that utility is increasing in quality for everyone, regardless of theta. And the high theta people have a higher willingness to pay than the low theta people, again, for any quality. That doesn't fit all products, but assume that there are these high types and low types and the high types are willing to pay more for everything, as if they're wealthier people with a lower marginal utility of income.

Now imagine that in equilibrium the firm is going to sell q units. Given that the high theta people have the highest willingness to pay, if you're going to sell-- let me put my-- theta is uniform 0, 1 here. If I'm going to sell q units, what's going to have to happen in equilibrium is everyone whose type is above 1 minus q buys the good, and anyone whose type is below 1 minus q doesn't buy the good. Because if someone here is willing to buy the good, anyone with a higher type is strictly willing to buy the good at the same price. So if you're going to sell the good, this is how it's going to look in characteristic space.

So therefore, what's your maximization problem? I'm going to maximize. I'm going to set it over, maximize over the quantity that I'm going to sell and the quality. If I'm going to sell to everyone in this interval, it's got to be that this person is just indifferent between buying and not buying. This person is willing to pay v of s 1 minus q, because 1 minus q is their type. So if I'm going to sell to these people and not to these people, it must be that this is the price I'm charging. So if I'm going to sell q units, my profit is q times v of s1 minus q minus c of s.

I now want to differentiate that with respect to s to say, what is the optimal quality to choose? Notice the q. This c drops out of the first order condition. So I just get dv / ds evaluated at the monopoly level of s and 1 minus q equals dc / ds. So what this says is that the effect on the unit cost of increasing quality is exactly what the effect of that quality increase is on the evaluation of the marginal consumer.

So the monopolist is choosing optimally-- choosing quality optimally from the perspective of the marginal consumer. And that's because the marginal consumer, if you raise the quality, that person will be willing to pay more, you can charge them more. And so you're improving quality if and only if the amount that the marginal consumers are willing to pay for that improvement in quality matches the cost of implementing that improvement in quality. So that in some sense, it's a partial optimality theorem.

What would a social planner do? Well, a social planner would be choosing-- so let me say, so this would be what the monopolist ended up doing. What would the social planner end up doing? It's maximizing over q and s. Integral over all types from 1 minus q to 1, v of s theta minus c of sd theta. Maximize that. I differentiate with respect to s under the integral sign. So I get the integral from 1 minus q to 1 of dv / ds minus dc / ds d theta.

The integral from 1 minus q to 1 of dc / ds is just that's a constant function. So it's just q times dc / ds. So I've divided both sides by q and I get this expression. The effect on the marginal cost of increasing quality is equal to-- this is just an average across consumers of dv / ds. So what the monopolist is doing is instead of equating the cost of the quality improvement to the benefit for the marginal consumer, it's equating to the benefit for the average consumer.

So there are two reasons in which the monopolist is not doing what we want socially. First is that monopolist is focusing on the marginal consumer rather than the average consumer. But the other thing the monopolist is doing is typically in this case, the monopolist is choosing this quantity. The competitive quantity that the social planner might choose might be down here. The social planner is going to serve-- for the social planner, there would be many more iPhones sold than there are with Apple as an iPhone monopolist.

And so notice that these effects often go in opposite directions. One is that the monopolist is focusing on the marginal consumer rather than the average consumer. Because you're focusing on the marginal consumer, you may make the quality too low. It may be that the average consumer wishes that their iPhone had a camera that was as good as the camera on the Google Pixel, but it doesn't have a camera that's that good because the marginal consumer is not willing to pay as much for the camera improvements.

So you make the camera lower than the average iPhone consumer wishes-- worse than the iPhone camera-- average iPhone consumer wants the camera to be. But the second effect is because only wealthy people and graduate students have iPhones, because relatively wealthy people have iPhones, all these people are not actually getting iPhones at all. And so it could be that the average-- with a social planner, you'd have this many iPhones sold.

The average consumer would be somewhere down here. And serving the average consumer, the iPhone would be even worse than it is today. It's because Apple is only serving the high of the theta distribution. Their marginal consumer is higher than the marginal consumer would be if they were serving everyone with a $300 iPhone.

But one corollary of this analysis is if you have a homogeneous population model, which people often do not thinking about it just to simplify things, you make all the consumers identical. If all consumers are identical, two things happen. One is the marginal consumer is the average consumer, because marginal and average are the same when all consumers are identical.

The second is, in a homogeneous population model, you're always making an all or none sales decision. If everyone is identical, you either sell to everyone or you sell to no one. And typically you sell to everyone faced with those two choices.

So in a homogeneous population model, there's a strong optimal quality theorem. That is, if you analyze something that's a homogeneous population model, you assume the monopolist is maximizing its profits. You're going to assume it's choosing the socially optimal quality, because there's no scope for those two to be different. Questions on that? OK. And this optimal quality reasoning, this will come in a few different times later in the class.

Durable goods. So in many cases, people are selling durable goods that provide services over many periods. You're selling washing machines. You're selling cars. You're selling phones. Consumers are buying it today, and they're using that same product in multiple periods.

One intuition you'll have sometimes is that this is a potential advantage. Firms can use this durable goods nature, the fact that people buy it at different times to move down the demand curve and sell the goods first to the people who have high values. You have whatever new Xbox is coming out. There's some number of people who line up outside the store, sleeping on the curb to buy the Xbox. They do not get it at the best possible price. They all pay full price.

And then 14 months later at Christmas, that same Xbox is selling for much, much less. And maybe what Microsoft has done is they've walked down that demand curve, first selling it to the high value people, then waiting for the high value people to buy, then selling to the next highest value people and in some sense trying to discriminate between people based on when they want to buy the good.

While somewhat appealing when you first hear the idea, if you think about rational, forward looking consumers, a durable goods monopolist is competing with itself. And you're competing with your future self. And if consumers think your future self is going to lower the prices of the Xbox, they're going to be hesitant to buy it on the first day that it comes out, because they're going to think, why don't I just wait a month and get it when the price drops?

Or why don't I wait three months and get it when the price drops or whatever? And then the consumers who might have bought it six months out are going to think, why don't I wait till Christmas or whatever, where it's the day after Thanksgiving. I can get that thing at half price 7 months, 14 months after it comes out or whatever.

So first thing I want to do is just think about let's suppose you do have this model with commitment power and with the forward looking consumers. What would you do? So simplest model of commitment power. I'm often going to do this-- I know in some subjects, especially macro these days, everything has to be an infinite horizon dynamic programming model.

There's certainly plenty of infinite horizon dynamic programming models in IO. Often, though, in class, I'm just going to use two period versions, because I think the two period versions give you a lot of the intuition with much less complication than the Bellman equations would.

So my simple model of commitment power. There are just two periods. Consumers can buy the good at either t equals 1 or t equals 2. So think about it. There's t equals 1, there's t equals 2, then the world ends and everything goes away.

Good is perfectly durable. Consumers, again, are going to have types theta distributed on 0, 1. Here I'm going to make theta the actual value. So then I'm going to want theta to have a distribution capital f of theta. And so the consumer's utility is going to be 2 theta minus p if they buy it at t equals 1, theta minus p, if they buy it at t equals 2, and 0 if they don't buy. So the reasoning is like a type theta consumer gets utility theta every period they have the good. And so they buy it at t equals 1. They get to enjoy the Xbox in two periods. If they buy it at t equals 2, they only get to enjoy it for one period.

I'm going to have a constant marginal cost of producing the good, and I'm going to assume the monopolist has commitment power and can set the price as p1 and p2 that it's going to charge in the two periods.

Basic observation in this model is that in this model, as I've written it, you don't want to walk down the demand curve. In this model, what you want to do is sell the good at the monopoly price at t equals 1 and commit to make no sales at all at t equals 2.

So in this model, and obviously there are some firms that fit this. If you look at Apple's iPhone pricing, it seems like that iPhone comes out and it's $999 and everyone knows that the iPhone 14 is coming out literally next week and the iPhone 13 price still has not dropped. They've set it at the monopoly price for iPhone 13s and they're going to keep it there and never sell another one until the 14 comes out.

Why does this happen? Well, I'm going to think of my theta line. So suppose we've got our consumers. They're arranged by theta. What's going to happen in equilibrium? If I have some policy that involves selling the goods in both periods, I'm going to have some high types who buy at t equals 1, some intermediate types who buy at t equals 2, and some lower types who don't buy.

Why is that what's going to happen? Again, it's a revealed preference argument. If this person is willing to buy the phone at t equals 2, this person gets a higher value. They strictly prefer buying it at t equals 2. So if anyone is buying it, everyone higher is buying it. So the people buying it have to be above everyone who never buys it.

And then buying at t equals 1. If you buy at t equals 1, you may pay a high price for getting the Xbox on the first day or you can wait and get it later at a lower price. Again, if you prefer buying it at t equals 1, then anyone who's strictly higher value strictly prefers getting a t equals 1 relative to t equals 2 because they're getting the extra period of enjoyment of the good. That extra period of enjoyment is increasing in theta.

So if I do want-- often in IO, it often makes sense to think about the choices not in price space, but in this data space. So if I want to implement this policy, how do I do it? Well, for this to be the cut off between people who buy at t equals 2 and don't buy, obviously you're going to have to have p2 equals theta 2 hat. In the second period, you're selling it at a price of theta 2 hat. So the people who have utility values above theta 2 hat buy at the second period. The people who don't-- utilities below that don't buy.

What about in the first period? In the first period, you can't sell to this guy at a price of 2 theta 1 hat, because he'd be getting 0 utility buying at a price of 2 theta 1 hat. He would just wait until the second period and buy it at this low price and get positive utility by buying in the second period.

So what you can do, the best you can do is to chart, set a price so that the utility he gets from buying in the first period exactly matches the utility he could get from waiting and buying it in the second period. And so what it tells you is that p1 is going to be theta 1 hat bigger than p2. So p1 is going to be theta 1 hat plus theta 2 hat, which would make sense. He's getting the one extra utility of a period of enjoyment. So you'd think about what's his incremental benefit from buying it early versus late? It's theta 1 hat is his incremental utility. That's what you're charging him.

So monopolist problem looks like this. Maximize over theta 1 constrained by theta 1 hat at least theta 2 hat. The price minus cost at t equals 1 times the number of units you sell at t equals 1, which is 1 minus capital f of theta 1 hat, plus profits you get from second period sales times the mass of people who buy the good in the second period, which, again, is the number of people with types between theta 2 hat and theta 1 hat. So that's what my profit function looks like.

Notice that I've got two different terms that involve theta 2 hat minus c. It's here and it's here. So if I group those together, it becomes my-- my profit function looks like theta 1 hat times 1 minus f of theta 1 hat plus theta 2 hat minus c times this plus this. And then these things cancel. And so my profit function is theta 1 hat times 1 minus f of theta 1 hat plus theta 2 hat minus c times 1 minus f of theta 2 hat. That's my profit function.

Let's assume that the constraint was not binding. So if the constraint is not binding, I actually want to walk down my demand curve and set a high price and then set a lower price in the second period. This is just to separate maximization problems. I maximize this with respect to theta 1 hat. I maximize this with respect to theta 2 hat.

Well, now let me go back to [14.]121 for those of you who have taken it. [14.]121 last year, [? Parad ?] did a lot of this functions with increasing differences stuff. If you look at this function theta minus c, 1 minus f of theta, this is like the monopoly profit function. This thing has increasing differences in c and theta, which means that this is a very general theorem. When a monopolist marginal cost is higher, it sets a higher price.

This is basically the monopoly pricing problem with a cost of 0. This is a monopoly pricing problem with a cost of c. So the maximizer of this is larger than the maximizer of this. So if I ignore the constraint, this is the monopoly price when your cost is c. This is the monopoly price when your cost is 0. This tells you that theta 2 hat would be bigger than theta 1 hat. So if I assume this constraint is non-binding, I get the contradiction that if that constraint is not binding, I find that theta 2 hat is strictly greater than theta 1 hat, and that contradicts the assumption that the constraint is not binding.

Why is this? In some sense, what's going on here is-- almost imagine with this durable good, instead of selling those Xboxes in the first period, I rent them to people. I get all those people waiting in line. I give them a one period rental on the Xbox, and then I take them all back and then I rent them again in the second period.

Anyone that I'm selling it to in the first-- all these people who get it in the first period, they're going to rent it back from me in the second period at a price of theta 2 hat. When I'm selling extra units in the first period, any extra units I sell in the first period were units that I was going to sell in the second period anyway. So the cost of selling them one period early is 0. I'm already producing them to sell in the second period. So given that I'm producing them to sell in the second period, there's 0 cost to me of selling them in the first period instead.

So it's as if those units I'm going to sell anyway have a marginal cost of 0 instead of a marginal cost of c. And so because I'm going to produce them and sell them anyway, they have no cost. So that makes me even more want to sell them rather than less want to sell them. So this doesn't make any sense.

So this would say is that monopolist selling durable goods, this whole idea, it's a strictly bad idea to walk down your demand curve, because you're taking these goods that you're going to produce anyway. They really have no cost to selling to people early, and you're losing those people's enjoyment. It doesn't make any sense. These ones, you have to be hesitant to produce because they have a cost. These ones have no cost of production.

So in some sense, someone at Apple saw this lecture, they realized this is what you do. You do not bring out the new iPhone and then make its price go down and let people wait for the price cut. You just bring it out. You fix the price forever. And that's the optimal monopoly pricing policy with durable goods.

Simple numerical example there. I don't need to go over it. So where would price drops come from in real life? Obviously we need to do something different from this model to get prices of new goods to come out. I think we do often have the view that there are many new things. The new product comes out. It's got a really high price. And then if you wait for it, the price will decline over time. Why does that happen?

One thing you can get is different preferences. Suppose the people get utility not from enjoying the good, for instance, but from having the Xbox before any of their friends do. I'm the one guy who has the brand new Xbox. Everyone has to come to my house to play it. I somehow get status in my friend group out of being the guy with the new Xbox.

Fashion clothing. Clothes come out. Fashion designers introduce it. Someone has a copy some number of months later. But the people who wear the $3,000 dresses get utility from knowing that they have the genuine $3,000 dress that no one else has because it's just the new good that comes out.

So if you give people preferences for owning goods that other people don't have, then it can make total sense to have some set of people-- in some sense, almost that's going back to an optimal quality theorem. If there are people out there who have a high willingness to pay for things that other people don't have, you create a high quality product, which is a product that other people don't have, and sell it to them. And then maybe once you've sold it to people who have those preferences, then you can sell it to the other people later. So that would be one way you can get early adopters.

I think another thing that often comes up in the tech sector is just costs that are changing over time or costs that are changing over time because of technological improvements, because of learning by doing, because of capacity constraints at the factories. Early on, it's very costly to produce goods. The costs are high. Therefore, you set a high price. As the technology level changes, it becomes cheaper and cheaper to produce that good. You get better at producing it. Your costs go down. You drop your price as a result.

Third explanation is firms may have a commitment problem. That is, if you think about how is this pricing supposed to go. I bring out my new phone. I price it at $1,000. Everyone who wants that new phone comes out and buys it at $1,000. And then I'm supposed to sell, in this example, without phones that break, I sell none for the next two years or whatever until my new product comes out. So I sell a whole bunch of phones. I keep the price high. Everyone bought it at t equals 1. No one buys it at t equals 2.

If you imagine the firm sitting there at t equals 2, they must be thinking once everyone has already bought the product, if this was not like any repeated game, imagine this was the last iPhone that was ever going to exist. Why keep selling 0? Why not take all those people who have low values? The other people are going to be disappointed when you cut the price and wish they had bought it later, but it's too late for them to actually buy it later. Once you've already sold the phones, it is optimal for you then in the second period to go ahead and cut the price and sell to more people.

And so if you don't have some way to commit yourself not to lower the price after all the early adopters have bought, you're going to want to lower the price. And then once you're going to want to lower the price and equilibrium, you will lower the price and then you'll be in the position of everyone in the first period thinking you're going to lower the price and not buy it because the price is too high given what you're going to lower it to and so on. So firms could be lowering prices because they lack the ability to commit to a pricing structure.

So for instance, if you think about Apple versus some company that's a new startup that may never have a second phone, you could imagine that Apple is able to commit because they have some reputation. A new entrant is not able to commit, because no one knows if they're going to be around. We don't know that they're a firm that's not going to cut the price. And so it may be that that's a disadvantage that you get if you're not someone who is known to be there for a long time and known to have this reputation for high pricing.

I put on the slides, I don't want to take the time to do it. In some sense, I think that over time people in IO believe less and less in commitment problems and in firms being unable to exploit their monopoly power because they have these severe commitment problems. But anyway, I put on the slide a simple numerical example of imagine theta was uniform 0, 1. Cost was 0. d of p is 1 minus p. What is it that you do? And without worrying about the algebra, you can look at it afterwards. What you do is-- do I have it on the next slide? Yeah.

Whereas in this model, if you were a monopolist with d of p as 1 minus p and two periods, the monopoly price is 1/2 per period. So you'd set a price of 1. And so what you get, instead of setting a price of one and then sticking to it and selling to half the consumers, the solution turns out to be a set price of 9/10 in the first period and then drop it to 3/10 in the first period. You sell to some people, but less than half in the first period, even though your price is lower. And then you sell to some people in the second period.

So two things that come out in this simple algebraic example is you get a decreasing price sequence. This is the commitment problem. The price sequence has to go down because if your price didn't go down, no one would buy in the second period and then the second period, you would be getting 0 profits instead of positive profits. So in a problem with commitment, you always get p2 less than 1/2 p1.

And you always get profits are lower without commitment than with commitment, because if the profits were better with this 9/10, 3/10 thing, that's what you would have chosen when you had commitment power. So commitment power is always the best you can do when you have commitment power. So what you do in these sort of problems with a self control problem is always you earn lower profits and you do walk down your demand curve. So firms lacking this ability to commit is one explanation. And firms are worse off because they can't commit to doing what they would like to do.

OK. The Coase conjecture is even more extreme commitment problem. I think in the 1970s or whatever, this was something the Chicago School of Economics took seriously that if a monopolist has no market power and has to price at marginal cost because commitment power is so severe.

So the Coase conjecture is let's take the extreme version of this where I set a price at t equals 0. I set a price at delta t, which you think of as one day later, 2 delta t, 3 delta t, and so on. And assume that consumers have utility theta e to the minus rt if they purchase at time t. So consumers are discounting. Firms are setting their prices over and over and over again, and the consumers have this kind of discounted expected utility and are forward looking.

The Coase conjecture says that under some conditions, the monopolist profits go to 0 as delta t goes to 0. So if you make this severe commitment problem that every day I'm going to choose a new price, every day I'm going to choose a new price with no ability to commit to that price, the theorem is that actually the more often you can change your price, the more severe this commitment problem gets.

And in the limit as delta t goes to 0, what you do is basically you do almost like that. You start your price out closer to 1, but then it just drops to 0 and it goes to 0 almost instantly. And in the limit as delta t goes to 0, and in fact, your price goes to 0 instantly and the monopolist earns 0 profits at all.

Why does this happen? Again, it's the example, the argument is really there's nothing else that could happen in this model. So suppose I did this example where if I have commitment power, I just set the price at 1 and never change it. So if I put time on this axis and p of t on this axis, this is what the monopolist wants to do.

And when you have no-- when you had two periods, I think I showed you that this is some sense. That's what the monopolist does. It sets the price here and then drops the price down. If you have three periods, the monopolist is going to do something like this, set the price here, set the price here, set the price here.

Imagine the prices didn't go to 0. If prices didn't go to 0, it must be when you do this thing and very rapidly, you get some smooth curve. Like this is what the monopolist does when it's changing prices every second. It just drops its price smoothly over time. But dropping your price smoothly over time at a finite slope just isn't going to ever be optimal. And why is that?

The reason is that imagine this was what you were going to do and you decided this time to instead of charging this price, I'm just going to jump ahead one period and charge this price instead and then follow the same path. So I just say, OK, this one period, I'm just going to jump ahead and charge that price instead of this price. If I jump ahead one period, I move all future sales back by delta t. The gain for moving those things ahead is first order in delta t. Because it's got an e to the minus delta t in it. I just move sales forward one period.

What have I lost by dropping the price from this to this? Well, I could have earned this price on this tiny, tiny number of units I sell in that period. I now get this price on those units. And so my loss from having jumped forward is this thing times this thing. That's second order in delta t, because the number of products I'm selling in each delta t unit interval is like a delta t, proportional to delta t, and the price drops are proportional to delta t.

So the gain for moving things forward would be first order. The loss for moving things forward would be second order. So you would always want to move things forward. And so what must be the case is it's not the case that these things are both order delta t. It's got to be that this thing is falling at a nearly infinite rate so that the loss and gains both become first order.

So anyway, so the Coase conjecture says there are some conditions under which this happens. There's a big literature on the Coase conjecture. If you read the Fudenberg, Tirole *Game Theory* textbook, not the IO textbook, it's got pages and pages and pages proving versions of Coase theorem. I think an IO right now-- I should say it's Coase conjecture even though it's a theorem, whereas the Coase theorem is really a conjecture. But the Coase conjecture is true. It's a theorem. People have proved it under various assumptions.

Why isn't it true? There's more of a literature in IO on why it doesn't happen, and I don't think anyone believes it should happen these days. Why wouldn't it happen. First of all, firms don't change their price every day. It takes time to change your price. In particular, it takes time after you change your price for all consumers who are supposed to buy at that price to buy at that price.

So if you decide I'm going to sell my phone for-- my Xbox at $399 instead of $459, you're going to have to wait a month for all the $399 consumers to get around to coming in and buying it. So it's not like the consumer's buying instantly. It's also costly. You're not going to change your price every day.

Firms like Apple can have reputations for not buying their goods. You can also have inflows of new high willingness to pay customers. So even if Apple didn't have this reputation, there are people who are very high value iPhone users whose phones break. And then when their phones break, those new high theta people show up wanting to buy a new iPhone. So in some sense, there's always high theta people showing up in the market. And with high theta people showing up in the market, that gives you another reason to keep your prices high.

Consumers not being all rational and forward looking. Again, that argument that I gave you was all the consumers are anticipating the future price drops. If you have some consumers who don't realize the price is going to drop, then you're going to want to exploit the consumers who don't understand that the prices are going to drop. That's going to change that first order condition I gave you.

There are also many strategic reasons. So, for instance, you can rent the good rather than selling the good. Like in the old days, it used to be that AT&T, when they were a phone monopolist in the United States, didn't sell telephones. They only rented telephones. And you had to rent your phone from AT&T. They didn't offer sales as an option.

When Xerox first started selling copy machines, they didn't sell copy machines. They rented copy machines. If you rent them instead of selling them, you don't face this commitment problem, because just every month you're renting to your current set of consumers. You're setting your optimal price to keep renting it. That is the monopoly price.

The other thing you can do is a tricky thing that seems like it's favorable to your customers. You can say, I have this very customer friendly policy. If I ever sell this iPhone to any customer at a lower price than I sold it to you, I'm going to rebate you the difference automatically. It's going to come back to your credit card. It's a best price guarantee that you never have to worry about I bought my iPhone and got ripped and got ripped off and someone else paid less than me.

Again, I adopt that most favored customer policy. I commit to it. Then once I've sold my phone to all the high value customers, I don't want to cut the price, because if I cut the price, I'm cutting the price for everyone. That's back to the same thing as the rental problem.

And firms can also strategically, I can just build a factory that only produces so many phones per month. I build a factory that only produces enough phones to sell to everyone at one size. I'm not tempted to raise my price.

So I think the general view in IO is that there are enough reasons why the consumer is not being rational. The inflows of high willingness to pay customers are such that the extreme Coase thing shouldn't happen. And if you were in a situation where the extreme Coase problem were there, firms would come up with a way around it. Obviously, there are some antitrust concerns with some of these, but there are ways to avoid the Coase conjecture. I think we don't necessarily think that this is so important.

And I should say that this is something you'll get later in the class. We start doing-- a lot of people in IO are doing structural models of dynamic situations where you're always assuming firms follow these Markov perfect strategies. This Markov maximization is, in some sense, always a firm with commitment problems where it's every period the firm is choosing the optimal action given that state. Our skepticism of the Coase conjecture, in some sense, should translate to skepticism of some of those models. Or at least think about do we really think the firms lack commitment power and are always going to be maximizing period by period what's optimal for them?

So then the final topic I want to talk about today, and I'll see if I have time to finish it, this is a more recent topic in IO. And if you wanted to do IO theory on monopoly pricing as a thesis topic, this would be the active area of IO monopoly pricing you could work on today. So I'll give an overview of information design. And obviously, I would say if you really want to study information design, Steven and Ian in the [14.]281 class do a lot of information design, and you would want to know information design at the level at which they teach it. But I think basics for this I'm going to cover here as well.

So motivating lecture for this is that it is an interesting observation that when you look at goods, there are some goods for which you seem to get an awful lot of information about what you're buying and some for which you seem to get very little information about what you're buying.

So an example of a good where you seem to get a lot of information is you want to buy a new laptop computer. You can get on the site and the firms that are selling you laptop computers, they're going to tell you exactly how many pixels are in the screen. They're going to tell you what processor is in it, what the processor speed is, what the battery life is, what the battery power is, how much the laptop weighs, what it's made out of. You can get a tremendous amount of information about every laptop that you're considering buying.

I don't know how many of you have moved here recently and tried to buy a mattress in the United States. You try to buy a mattress in the United States and there are lots of these mattress stores and they sell really expensive. It's often like this mattress is $1,200 or this one is $300. And it's just impossible to figure out what makes mattress A different from mattress B.

And you even think, OK, I'll go to four different stores, all of whom sell Sealy Posturepedic mattresses and comparison shop. And you'll see that this store has the Sealy Posturepedic Ruby and Diamond and Emerald, and this store has the Sealy Posturepedic Gold and Platinum and Bronze, and this one has this Sealy Posturepedic Cloud and Mountain and Valley or something.

Just every different store, even if it's the same brand of mattress, they have all different names on them and no one's giving you any attribute. The coils in this one are 0.73 mils thick and the memory foam has this compression per unit thing. And given that you can't actually sleep on them in the store, it just seems like there's no information whatsoever to compare these matches and figure out do I want to pay $100 for this one instead of that one.

Wholesale used cars are almost one of the most interesting examples of explicit lack of information. So you turn your car into a dealer. Let's suppose that you had a car. I have my Honda Civic. I decide I want to buy a Toyota. I turn in my old Honda Civic to the Toyota dealer when I buy my Toyota. And then my Toyota dealer doesn't want to have a Honda. So he takes and sells it at the wholesale level to some Honda dealer somewhere who wants to have a used Honda on their lot.

You go to see these wholesale used car auctions and it's a remarkable thing. They have these giant parking lots in the middle of nowhere with some big building next to them. All the used car dealers show up with cars to sell and have their cars listed for sale. And then a bunch of people have the cars listed to buy.

And the rules for the wholesale used car auction is they drive the 2013 Honda Civic, 63,000 miles, up there. All the buyers are allowed to look at the car from the inside. A guy who works at the auction house starts the engine. You can hear that the car runs. You can look at it from the outside.

You may not open the hood and look under the hood. You may not look under it with the mirror and see whether the muffler is rusted out. You may not have your mechanic try out all the parts. The car is wheeled up there. Everything stays closed except for the guy from the dealer. He turns it on, runs the thing. They then auction off in 45 seconds and move on to the next car.

So basically, they provide no information whatsoever on the car that's being sold other than Honda Civic 2013, 6,300 miles. Here's what the engine turns on. Here's what it sounds like. This is what it looks like from the outside. And then the next one is sold. And it's clearly an intentional design to provide zero information to the buyers.

These examples, these are ones where we think of as these are clearly pure monopoly pricing decisions that somehow the mattress firm, the laptop firm, the used car firms, they've decided that providing this information to their consumers is what maximizes their profits, because they're controlling the information flow.

There are other examples, now, I guess, which has [INAUDIBLE] another side, which is that there are now plenty of intermediaries that make money by providing information about products and helping consumers find products they want. So, for instance, Yelp helps people find restaurants.

Amazon sells many things that are not Amazon products, helping you to find what's the best phone for me to buy? What's the best tent for me to buy? What's the best shirt for me to buy? What are the best pair of hiking boots? Whatever. Amazon is selling these products that they don't make and just providing you information about them.

One concern that you can think about, what does Amazon want you to do? Amazon wants you to buy products from Amazon. Amazon therefore wants to provide you with consumer surplus when you're buying these third party products, that you buy the third party products, you get consumer surplus. You keep coming back to Amazon as a place to start your shopping trips.

What information problem they have is that let's suppose they provide you with some number of information about the good. If they provide you with information that lets you figure out, yes, this is the perfect pair of hiking boots for me. Then the hiking boot company knowing that Amazon is selling it only to the consumers for whom it's the perfect product will set a really high price and the consumer surplus will be low.

So in some sense, Amazon has this problem of if they provide information that lets you figure out exactly how good the product is for you and what is, in some sense, trying to maximize the profit maximizing set of information, then it may be the monopolist then just raises the price and extracts all that surplus and then the consumers get no surplus. And then Amazon isn't the place to start shopping, because it's the place where, yes, you found what you wanted, but yes, you also paid a super premium price, and therefore you didn't get any surplus anyway.

So these kinds of-- I think it's the platform application and the developments in theory of this growth of what we call information design. Thinking about how would a firm that controls information about products want to provide that information either to maximize its profit, selling them directly, or to maximize the consumer surplus of the consumers who it's not selling them-- who are-- if it's an intermediary, it wants to maximize the consumer surplus rather than the profit, or it might want to do something in between.

So to think about an example for illustration, let's assume that we have a new product. Maybe this new product is Clover has a new fried cauliflower sandwich with some kind of Chinese sauce on it or Indo-Chinese fried cauliflower sandwich, which I think they do have. So suppose that fully informed consumers, if they really understood what the sandwich tasted like, would have values v distributed uniform 0, 1. Or maybe in this case, uniform 0, $18. Suppose that the monopolist has a constant marginal cost of 1/3. So $6 in this example.

What are you going to do if consumers are fully informed? They have values uniform 0, 1. You have a cost of 1/3. You maximize p minus 1/3 times 1 minus p. Monopoly price is 2/3. So we get this situation of your cost is 1/3. You set a price of 2/3. 1/3 of the people buy. So if I multiply by 18, I have a cost of $6. I set the price at $12. 1/3 of the people buy my sandwich. I make $6 each. I make $2 in profit.

So that's clearly one thing that Clover could do is just tell everybody everything about the sandwich, somehow convey to them what it tastes like. I don't know quite how they do that, but they convey it. People get these values. Maybe they give free sandwiches to everyone, so they all learn their values. They learn their v's and then they monopoly price relative to those values.

What else could you do? Imagine that consumers learn their values. Consumers have to learn their values somewhere. I know there's a new Clover sandwich. How would I know whether I like it or not? I know that I'm rational, so I know that v is distributed uniform 0, 18. But I don't know whether I'm one of the people who's going to this sandwich or one who's not going to the sandwich.

Imagine one thing Clover could do is provide no information whatsoever about the sandwich. Don't even tell you it's got cauliflower. Don't even tell you what the spice is. Say, this is the new Clover sandwich. If you're a rational, you're going to know that, OK, there's a new sandwich out there. My value must be somewhere between 0 and 18 for that sandwich. I don't know what that value is. Or uniform 0, 1. So here I would think on average that sandwich is worth $9 to me not knowing. It's just in a brown paper bag. They've kept. It's a secret sandwich. My value is 1/2 or 9 in that example.

So what the monopolist can do is sell these sandwiches in the brown paper bags to everyone at a price of 9. And so instead of selling it to the-- I had the previous example, I'd sell it to 1/3 of the people at a price of 12 and earn 1/3 times 12 minus 6 is $2 profit. Now everyone is willing to buy the sandwich in the paper bag. They're only willing to pay $9 for it because they don't know whether they're someone who likes the sandwich or someone who hates the sandwich. But I get 1/3 times 9 minus 6 is 3. So I've made more money by hiding what my product is.

And how I made more money by hiding what my product is, well, social surplus has gone down. There are all these people with low values for the sandwich, values below 6. That sandwich should not be sold to them. There's this inefficient production that there are people who don't even like the sandwich it's being produced and sold to. But I've eliminated the monopoly deadweight loss that came from these people with values between 6 and 12 not getting the sandwich because monopoly pricing at 12.

And so because I'm extracting all the consumer surplus, I've driven consumer surplus to 0, total social surplus has gone down. But because consumer surplus went down even faster than social surplus went up, suppressing information is transferred surplus and the monopolist is better off. So this is just an extreme example where I hide information about my product, my profits go up, and it goes up because I've made consumers homogeneous. By making them homogeneous, I can extract all their surplus. And that's socially wasteful, but it's a transfer from the consumer to the monopolist and it makes the monopolist better off. So anyway, that's one thing Clover could do here.

Second observation is that you can also have examples like this where providing less than full information actually makes the consumer surplus go up. So here's an example. Again, think about Amazon or Yelp trying to sell through Clover, provide information about Clover sandwiches to people, and give the consumers more surplus by using the platform.

So suppose what happens here is that-- so think of Yelp as the firm designing the information structure. And suppose that Yelp, because they've read all your reviews that you've ever written of restaurants in the past, they know how much you're going to like the Clover sandwich. They've got enough things of people who both ate this sandwich and had the same taste as you at all previous restaurants, that Yelp knows whether or not you're going to like the sandwich. Because in some sense, they implicitly know what it tastes like relative to your tastes.

And so suppose what Yelp does is it commits to this partial information policy. If your value is between 0 and 1/3 or 0 and $6, it tells you, you are someone who is going to have a low value for that sandwich. If your value is between 1/3 and 7/9, so that's between $6 and $14, it's going to you, you have a medium value for that sandwich. And if your value is between $14 and $18, it's going to tell you, you have a high value for that sandwich.

So it's like it's taken the sandwich, it's wrapped it in a paper bag and not told you what it is, but it says, based on what we know about you, your value for this sandwich is low. Your value for that sandwich is medium. Your value of that sandwich is high. And then it sends you into Clover with that knowledge of you only can buy from Clover through Yelp in this example. I mean, it's more plausible with Amazon. And so all you're provided with is the information of my value is low, my value is medium, or my value is high.

So what does Clover do when it's in this world? Well, the people who are told the value is low think that the value is 1/6 or 3. Those who think that their value is medium think their value is 5/9, halfway between 3/9 and 7/9. And those who are told the value is high in expectation think their utility is going to be 8/9.

So Clover is facing a strange population where we now have, instead of having consumers with values uniform on 0 to 18, we now just have 3 point masses. We have a point mass of 1/3 on people who think it's worth $3 to them. We have a point mass that's this big. 4/9 of the people who think the thing is worth 10. And then we have a smaller point mass 2/9 who think the thing is worth 16.

So what does Clover do if it's faced with this demand? There are some people who think it's worth 3, some who think it's worth 10, some who think it's worth 16. Well, obviously, you're going to price the good either at 10 or 16 in this example. Because everyone has a value of-- if I price a value more than 10, only these people buy, so I might as well charge 16. So I basically have two choices. I can price it at 10 and earn 10 to the minus 6 and sell it to 2/3 of the people. Or I can price it at 16 and get 16 minus 6 and then sell it to 2/9 of the people.

And I think the way I did these numbers, selling to both groups at 10 is better for them than selling to only the high types at 16. So what happens is I get them to sell the sandwich to both the medium and the high types. And if you look at what I've done, I've worked out the calculations here.

I've gotten them to sell the sandwich. Instead of having the monopoly price of 12, the monopoly price was sell it at 12 and everyone with a value above 12 buys. Here I've gotten them to sell the good at a price of 10 instead of 12. And so by getting them to-- by some sense, creating this extra mass of consumers who think it's worth 10, I've shifted down where the monopoly price is and gotten clover to sell the good at 10 instead of 12.

And it turns out, if you work out the algebra here, the profit of the monopolist is, I believe, lower. Nope. So actually, sorry, it's here. So consumer surplus is higher than it is under monopoly pricing. And actually, in this example, the profit of the monopolist is also higher. So this was the no information treatment. This was the full information treatment. This one I've actually bumped up consumer surplus and profits relative to the monopolist is done.

Anyway, so I am out of time now, so why don't I leave it there? I will do my final slides on information design next time. And I should say jumping ahead, as I said, I'm going to do theory on Wednesdays and empirical work on Mondays for the immediate future. So in Monday's class, I'm going to talk about two empirical papers on monopoly pricing, Chevalier-Goolsbee and Gentzkow-Shapiro.

So be useful for you to-- I will discuss what's in the papers, but probably useful for you to try to look at the papers beforehand, try to read what's in them, get some sense for reading IO papers. I will say that Gentzkow-Shapiro is a bit difficult paper to read if you haven't done a lot of econometrics. We are covering the demand estimation in week four of this course.

And in the first three weeks of this course, you may be reading papers on demand estimation that are harder to read. I would just in the first three weeks skim through them if you haven't done GMM estimation and things and have trouble reading those papers. OK. Thanks. Hope you enjoy the course, and I will see you on Monday.