[SQUEAKING] [RUSTLING] [CLICKING]

**GLENN ELLISON:** Today I'm going to do some empirical literature on entry. And I wanted to start with this really classic paper by Bresnahan and Reiss. And Bresnahan and Reiss-- it's actually-- it's an entry paper, but it's really a static competition paper.

So the research question here is, how empirically does level of markups and competition change when you have more firms in the market? Is it true that-- I guess maybe it is the entry graph. What does this graph look like when-- what does the price/cost margin look like as the number of firms-- or their profits look like as the number of firms enters.

And Bresnahan and Reiss are starting from this classic NEIO approach of thinking that it's very hard to estimate profits because it's very hard to estimate costs. In many businesses, there are lots of intangible costs. And you can't just [INAUDIBLE] to get that. So we want to use something that's readily observable, and the idea is that entry is very clearly observable, whereas the [INAUDIBLE] costs observed are not.

The other idea in Bresnahan and Reiss is that in IO, one of the challenges is always, how do you get a cross enough-- a big enough cross section of observations to run the significant statistical estimates? And Bresnahan and Reiss had this idea of looking at isolated small towns.

And so if you look at the United States, one thing that's nice about the United States is it's very, very big. And there are many parts of it where there aren't a lot of people. And so that gives you a lot of towns where there are towns, and then there's very little around them.

And so if you think about-- you go search around the United States, it's a big enough country that you can-- if you want towns-- and in Bresnahan and Reiss's case, what they're looking for is-- I forget what the cut-offs are. You have towns that have at least-- you have cities that are isolated in the sense of-- I think they're all 100,000 miles from any city of 100,000-- they're 100 miles from any city of 100,000 people. And they are at least 20 miles from any city of 1,000 people. So these are all towns that are just completely isolated from anything else that you would-- any other population centers.

And they note that there are a lot of them-- and, in fact, I think they-- Peter tells stories of driving around and wanting to verify the data by just driving around to all of these small towns and then asking them, where can I get a plumber in this town? Where can I get-- where can I buy tires for my car or whatever?

And so what they did is-- most of the data collection was done through business directories. They basically just looked at a US map, found places that were completely isolated. And then we used to have this thing called yellow pages, which was like a phone directory, where you would look up phone numbers in a book.

And so people would use these books to look up phone numbers. They would find out who was advertising as a plumber in each of these towns. They also use just business directories. And then they visited them to see, was this data being accurately collected?

The paper-- it's a single cross-section. They don't observe firms entering in or exiting. But what they're going to infer is that sometime prior to the point at which they're collecting the data, this firm's-- classic game is firms choose in/out. And then firms choose prices or quantity star of n-- q star of n.

They paid a fixed cost. They assume that this game has been played at some point in the past. And so what we're observing in each market is, what's the equilibrium number of firms in that market?

The basic thing the paper then tries to do is say, if we can observe the number of firms in a large cross-section of markets, that's going to give us a means to estimate profit levels without any data on prices or on costs or on demand.

And so consider of basic model. You have the entry decision. You have a firm that has fixed costs, F. There's going to be symmetric competition. So if there are N firms in a market with S people, you're going to get S over N-- so the total number of customers divided by the number of firms times pN minus c D of pN. And so they're going to call that's S over N times pi N.

So pi N is the per population profit-- or per population served profit in this market. It reflects both sales you make to each customer-- there's obviously only some fraction of the consumers that are buying from you. And so this is profit per person that's in your part of the market.

So if you have a town, and roughly, it's divided like this, and each firm gets one quarter of the market, it's your per customer profit for the market that you serve.

We're going to have N firms in equilibrium if S over N times pi N is bigger than the fixed cost. But S over N plus 1 times pi N plus 1 would be less than the fixed cost. And obviously, this is decreasing-- this is smaller than this, both because you're dividing the city into more fractions, and the per customer profit goes down because the price goes down.

And what we're going to say is what asset would be the minimum size in which-- so, for instance, let's put N on this axis, and then you can put the number of doctors in a town on this axis. And what you would observe, if you could draw this graph, is towns that have no-- that are very small don't have a doctor.

And then there's some threshold-- maybe it's like 800 people-- where now there's a doctor in the town. And then some other threshold where 2,000 people were. Now you have to have two doctors, and then some other threshold where you tend to have three doctors-- 2, 3.

And so the idea is that we can estimate these cut-off-- we can estimate these curves. Obviously, in practice, you're going to see a lot of zeros. You may see some doctors even in very small towns.

You'll still see some towns without doctors here, but you'll see a lot of 1s, and maybe you'll even see some 2s. But there's going to be, in general, some curve like this where this is-- typically, you need to have S1 customers to have one doctor. You need S2 to have two doctors. Now, here, you need S3 to have three doctors.

And so what they want to do is estimate what does this look like? How many people do you need to have before the first doctor comes in? How many do you need to have before the second doctor comes in? How many do you need to have before the third doctor comes in?

Possibly you can estimate this just by plotting the cities with the population on the x-axis, the number of doctors on the y-axis. Fit a curve that looks like this through them-- see where the jump points are.

So if you assume that the fixed costs don't vary with N, then SN over N is going to be pi N times pi N is going to be F. And also, SN plus 1 over N plus 1, pi N plus 1 is going to be F.

So both at this point-- at this point you know the second doctor is just indifferent to entering or not because if it was smaller, he wouldn't enter. If it was bigger, he would enter. Here, where the third doctor is in different entering or not. So both of these, the profit that you get if you do come in, is exactly F.

And so if you just set this equal to this, and then you divide the-- bring the pi N down here in the denominator, you get that the profit-- with N plus 1 firms divided by the profit of N firms is N plus 1 over N, SN over SN plus 1. You could also write that as pi N plus 1 over pi N equals SN over N divided by SN plus 1 over N plus 1.

So the idea here is that imagine it takes 1,000 people in a market to have one doctor, and S2 is 4,000. Then this is-- S1 over 1 is 1,000. You need 1,000 patients in order to be worthwhile to open a doctor's office.

But then here, S2 over 2 is 2,000. So that would say you need 2,000 people to make it work-- you need 2,000 patients for the new doctor to make it worthwhile to have a second doctor's office. And from that, we infer that it must be that pi 2 is just 1/2 of pi 1. So this ratio of this to this is how much profits per person must have declined for the second doctor to not come in until the second doctor is going to get many more patients than the first doctor got.

And so anyway-- so that's the really powerful idea of this paper. We don't have to have any data at all on prices or quantities. And yet, we can tell how the ratio of profits declines as more firms enter the market by just looking at how many more customers have to be there for the second firm to want to come in. Question about that basic idea? OK.

The previous slide suggested what Bresnahan and Reiss were going to do is put population on the x-axis, number of doctors on the y-axis, and then just draw a curve like this that best fit the data. That's not exactly what they do because they want to have a lot of covariates in the regression that are also affecting all of the elements.

So basically, they set up the models. First, they have something they call the effect of population. The effect of population in city I is its actual population plus some X-- some explanatory variables, X, times some coefficients, gamma.

Those X's, they use for things like the rural population within 5 miles of the city, the rural population within 10 miles of the city, the growth rate of the population. I forget what else goes in there. But they have the actual population of the city plus some controls for nearby population.

Per consumer profits-- again, the idea is that you would have per consumer profits, alpha 1 under a monopoly. Profits alpha 1 under the monopoly. Alpha 2 would be duopoly profits. Alpha 3 would be 3 firms.

But then you're going to have other variables that affect the profitability of a market. Those variables can include the per capita income in the city, the age distribution. Maybe higher income cities you earn more profits per customer or things like that.

The fixed costs-- again, it's mostly a fixed cost. They do allow there to be covariates that affect the fixed costs across cities. So, for instance, cities that have more valuable agricultural land, it may be more expensive to open up a tire dealership than in places with cheaper agricultural land. Because the land that you need to buy to set up your tire shop costs more. So they have a fixed cost to also vary with some covariates.

And then the basic thought is that-- so I have all these elements. I've redefined Si, pi iN, and Fi to be their covariates. And then my model is if there were N firms in the market, the profits would be Si over N, pi iN minus Fi plus epsilon i, where epsilon i is just a normally distributed random variable.

And so in this model, when are we going to observe in firms? We're going to observe in firms if Si N-- sorry-- Si over N minus Fi plus epsilon i is greater than 0. And Si over n plus 1, pi iN plus 1 minus Fi plus epsilon i equals 0.

So we have these two things that have to both be true for there to be exactly N firms in the market. But you could write the first one as epsilon is bigger than something, and the second one is epsilon is less than something. So if you view epsilon as a random variable here, a random shock to fixed costs, we're going to observe N firms whenever epsilon is in this interval. It's between this and this.

And so given the param-- given all the data and the parameters, you can compute the likelihood of this as just the normal CDF evaluated at this point minus the normal CDF evaluated at this point. And so then you can estimate all the parameters-- gamma, beta, alpha, delta, gamma-- this one was lambda. You can estimate all of those by maximum likelihood. And that's the basic estimation technique in the paper.

People-- I know we still teach ordered probit models in econometrics. This is what people refer to as ordered probit model, where you have discrete data produced by the discretization of something that has a normal error in it.

And obviously, the main things of interest here coming out of this estimation are alpha 1-- these were the objects of interest. What is the per customer profits under monopoly? What are the per customer profits under monopoly, duopoly, whatever?

But somehow in the paper, they really liked the idea of the S's as the thresholds. And so they basically estimate these things. And then they convert from all the-- they convert from all the parameters back to what is the mean S1, mean S2, mean S3, holding all the covariates at their sample means. And then they do a lot of the presentation just talking about what are the S's?

OK. Anyway, so data samples-- as I said, they've got data on a whole bunch of businesses gathered from 202 small towns in Western US. These are small places. The mean population is 3,740. And, as I said, they're all at least 20 miles away from any place that has 1,000 people in it. I went to college with a guy from Wyoming who said he was 90 miles from the nearest movie theater. So there are parts of the United States that are like that.

Here is the data sample showing the full set of occupation-- or businesses they started with and the number of firms they get. They end up focusing on the ones on the top part of the table. There are some paper-- some results in papers about things farther down in the paper.

But what they wanted-- in some sense, if you want to talk about how the competition changes with the number of firms, you want there to be a significant number of cities that have zero firms, 1, 2, 3, 4, or whatever. And a lot of these businesses, it turns out, you never get-- a lot of these businesses are apparently too big businesses to have many of them in the towns that they're working with. So they don't get much beyond N equals 2.

But if you look at the druggist-- I assume this means pharmacies-- they're very relatively rare to have no pharmacy. But then you do get 1, 2, 3, and some 4s and 5s and beyond. Same things with the doctors, dentists, plumbers, tire dealers is-- in some sense, they focus on these because they all have at least some number of markets with 3 and some number of markets with 4.

There was also a question of how well-defined are these markets? Like, the sum of these things, like the barbers or whatever-- people who are not barbers can cut people's hair. And so they're sort of-- it gets harder to count how many barbers or beauticians there are in a town because it's just somebody who works somewhere else-- also cuts hair. People know they can go to that person or whatever.

So it's hard to-- whereas, things like dentists, I guess, are easier to identify that someone is a dentist. You don't get just guys who come over and drill holes in your teeth on weekends.

So what do the entry thresholds look like? So, again, the numbers I put up on the board over there are-- most of these businesses tend to start to exist in towns, and they have fewer than 1,000 people.

So the doctors-- the cutoff is 880 people. The dentist-- the cut-off is 710 people. The druggist-- it's 530 people. The tire dealers, it's 490 people. You need to be bigger to have a plumber. That's over 1,000-- 1,400 people.

But then if you look at-- let's just look at-- the druggists is one of the most stark examples. It seems like you need 500 people to have one. You need 2,000 people to have two. So that's when you have 1,000 customers per drugstore. To have three, you need 5,000. So that's like 1,600 people per druggist.

And when they put the ratios over here and make it clear, the S2 over S1 ratio-- this is how many customers per firm you need to have to have-- so this is-- they use little sn to be capital SN over N, which is always tough for when you're putting things on slides that are small.

So anyway, these are the ratio of the per customer-- the per firm number of customers, which you can think of as this is the thing that's supposed to estimate pi N over pi N plus 1. So this is supposed to be pi 1 over pi 2.

And so what this is saying is that in many of these industries, profits are dropping almost in half when you have-- you have this when you have a second firm, or per customer profits dropping almost in half. It's like profits-- divide-- you divide it by either 1.9 or by 1.8. So that's like a 45%, 50% drop in profit when the second firm comes in.

The one exception appears to be plumbers, where we don't get the first plumber for a while. And so it seems that the second plumber is also serving 50. You get one plumber when you have 1,400 people for the plumber to serve, a second plumber when you have 1,500.

If you look at what happens when you go-- so one firm to two-- it's a big drop in profit. When you go from two firms to three, it varies by industry. But most of them, other than the plumbers, we have a 1 point-- a 10% drop, 30% drop, 60% drop or whatever. So it, again, is a-- or 1 over 1.6.

So, again, you get substantial drops often going from 2 to 3. By the time you're out here, though, the S5 over S4 numbers, here, it seems like competition-- whatever competition you're going to get is all done by the time you have four firms. Like four firms are getting you just as much competition as five firms is giving you.

So the conclusion here of how basics of competition seems to work is there's a big difference between monopoly and two firms. It seems like prices keep going down from 2 to 3. And then it seems like we asymptote pretty quickly to things are not getting any more competitive. Any questions?

I mean, I think it's a fantastic paper in that this idea of entry and telling us about profit levels and that being an optimality condition we can use has just been tremendously influential. As far as the applied results go, I would say the biggest difficulty with this about-- with regard to this as an applied result is this model is interpreting what we're seeing as the equilibrium N star of this game.

So the equilibrium N star of that game implies in part that everybody who is there is happy they're there. And everyone who's not there is happy they're not there. I think the happy they're not there part I think is very clear. The question is, what do we make of all these doctors in towns of 800 people?

Is the doctor in a town of 800 people because the doctor is happy serving 800 patients? Or is the doctor in a town of 800 people because it was a town of 1,500 people and a thriving farming town in the 1950s when the doctor moved there. And the doctor is now 75 years old and doesn't want to leave the patients the doctor has had for 40 years, even though the number of patients is dwindling relative to what it was when the doctor got there.

And if you think about it, if there's this sort of-- you pay this sunk cost to enter, and you don't get your sunk cost back when you leave, it could be that people enter-- the entry-- the things that we're seeing are not entry thresholds. They're exit-- we're seeing exit thresholds. And the exit thresholds might not be counting all fixed costs.

Another potential issue is just that counting firms is more difficult in their application than in others. How do you count someone who does some plumbing on the side, or someone who does some hair cutting?

So I think the ones in the licensed professions, like the dentists and doctors, are probably better with that. But the ones that are more like the businesses we normally think about price competition operating in, some of those have this issue of it's going to be hard to identify businesses, and ensure that we're [INAUDIBLE].

For instance, the result on plumbers could be that you don't get a lot of market power with one plumber because there's always somebody else in town who you can call who just owns wrenches and can come over and fix stuff for you, even if they're not officially plumbers. Questions? OK.

So then the next step I want to go to was Berry and Waldfogel. So this is very much following up on of theory of efficiency of entry I did last class. And so, as I said, there are reasons why entry is going to be higher or lower than we would like it to be.

And Berry and Waldfogel are trying to say, let's pick an industry. Let's say in this industry, how does it work out? Are there more firms than we would like or fewer firms? How big are the welfare losses from the number of firms that we observe?

The industry that they pick is commercial radio. And I think the argument is that there's-- you need a license to open a radio station. There are some places where you might-- where there's a-- you can't open a radio station and broadcast on any band because you'd be interfering with the stations right around you. But in most cities in the United States, there are bands that are available.

There's obviously a fixed cost of applying for the government license and opening up and just putting up a transmission tower. But that's fine. That's exactly all the Fs.

And so this is an example of an industry where we think there is a substantial F. Because to have a radio station, you need equipment. You need people there. But then we can think about how are the-- how is the prices affected.

Now, radio is free. So in some sense, the price you're used to thinking about for radio is zero. But radio is a platform market. So there are two types of customers at radio stations. There are listeners who typically pay zero to listen. But then there are advertisers who pay money to put commercials on.

And so the standard way that-- because the price is free, the standard way that the FTC or antitrust regulators tend to treat radio is to focus on the business of producing advertising. So they think of as a radio station is-- they have this technology where they buy rights to music. They play that music. That music produces ears of people. And then they sell those ears of people to advertisers who want to advertise to them.

So when we're thinking about price in this market, we should be thinking about price of advertisers-- price of-- think of the price as the price of a listener, and think of the firms as having a technology for producing listeners. So this technology for producing listeners involves something about radio music quality or content quality. And then that produces people that they then sell to the advertisers. So the advertisers are the consumers. The people are ignored.

And one thing you might notice, if that's your view of radio, you'd be like, wait. Wouldn't that mean like whenever two radio stations are proposing merging, what would be considered bad is if by merging the radio stations are going to raise the price of advertising and have fewer commercials per hour and play music instead.

And that actually tends to be how the US antitrust authorities have largely conceptualized radio. That is, if there's something going on that's going to result in fewer commercials per hour and more music or more news, then they're skeptical of it. If it's something that's going to result in even more commercials per hour, that's good because that's stopping the deadweight loss on the advertiser's side.

Obviously, Berry and Waldfogel mostly adopt that. They are clearly well aware that radio is also something people listen to, and they talk about that at the end when all the analysis is done, thinking about radio stations as things that produce advertising to sell to advertisers.

So we've got 135 US markets. We're trying to compare with what would be socially efficient and talk about the welfare losses. We're going to focus on the radio stations as firms that have advertisers as customers.

They're going to assume-- they're starting with the very first model of entry that I did. So they're assuming that the ears that are being produced are a homogeneous good. So every listener is worth exactly the same amount to every advertiser.

And, in practice, obviously, there are going to be stations playing music from the 1950s, where they have lots of ads for reverse mortgages. There are going to be stations playing much more up-to-date music that are advertising concert tickets.

Somehow they just they just abstract away from that. It's just-- they're just going to go to listeners are a homogeneous effect. There's just a downward-sloping demand curve for the more listeners you have, the price-- or the more listeners you-- the more commercial slots you have to sell, lowers the price per commercial slot. And that's how that market is going to work.

Anyway, so the-- what's also good here is they have purely fixed costs and no marginal costs of producing listeners. So you get even simpler welfare function than the one we talked about before. So the welfare-- if you have-- and radio stations, it's going to be the integral from 0.

So L star of N is going to be the number of listeners per station when you have N stations. So it's just the integral from 0 to N L star of N. P of x dx-- the willingness to pay for the very first-- for each listener integrated out over the total number of listeners.

So this is the advertiser willingness to pay for listeners. And this is the gross surplus, is just the area under that curve, minus the fixed costs that are incurred in generating all those listeners.

And so the things that we need to estimate are this per firm listening function L star of N. We need the advertiser inverse demand function, L of N. And we need the fixed cost F. So to evaluate welfare in this perspective, you just need to estimate those three objects.

How are we doing this? So the first [INAUDIBLE] of listening, this is where you get the sense of how much NEIO people like discrete choice models and discrete-- discrete choice models of estimating things because they don't have any price variation. They're not actually trying to estimate the consumer preferences. Despite not trying to estimate the consumer preferences, they're still using a model with consumers having discrete choice preferences, choosing among stations or the outside good to generate their listening function.

So this is-- people know what nested logit models are? So, OK. So anyway, this is a-- what we call a nested logit model. Do you remember I sort of-- we went with the standard logit model, the utility the consumer i gets from listening to station j is delta j, which is like the quality of station j plus epsilon ij, which is like the logit type.

In the nested logic model, you add a second shock. So the epsilon ij is the shock that's independent across stations. I like the music on this station. I don't like music on this station. That's additive to my preferences.

But then there's also a common shock, vi. And the vi is a random variable that's drawn randomly at the level of the customer, not at the station. And it is their taste for music as a whole versus not listening to the radio.

So we have these two shocks-- one common shock that customer i gets regardless of which station they listen to. And then one set of-- and then a set of idiosyncratic shocks that are independent across stations that you get if and only if you listen to that particular station.

And the idea here is by having a sigma that can vary between 0 and 1-- so sigma equals 0 is the pure logit model, where when you add another station, it draws customers equally from all of the other stations and from the nonlisteners.

And in radio, nonlisteners are most people. In most half hours, most people in the United States are not listening to the radio. It's actually surprisingly few. It's about-- I think it's like 85%. So typical half hour, 85% of people in United States are not listening to the radio.

But so then if you have the pure logit, that means when a new station attracts someone, it's 85% chance that they're expanding the market and choosing someone who wasn't listening to the radio to begin with.

The other extreme, sigma equals 1, this shock goes away. It's all the common shock. And when it's all the common shock, then it's 100% business dealing. If you only have common shocks, and the idiosyncratic means are essentially zero, you add a new station, it's just going to pull people away from the other stations. And it's only the people who have the vi's that are big enough to overcome delta that will listen to radio regardless.

Anyway, so this-- they have this model. It's got this free parameter sigma. And that's going to let them say how important business dealing is. And it's very important to have that because a business dealing-- we know that there's a lot of business dealing if there's too much entry. If there's very little business dealing, then we might have insufficient entry.

If you have a totally-- if this model was totally symmetric, if all the deltas were the same, if all the stations were exactly the same quality, then the nested logit model N l star of N would be N to the 1 minus sigma power over e to the minus delta plus the N to the 1 minus sigma.

So again, for sigma equals 1, total listenership is just completely fixed at 1 over 1 plus e to the minus delta, regardless of the number of consumers. But then when sigma goes to zero, this thing is like-- it's asymptoting towards 1, but it's going up initially linearly then. Because it's like N to the first over this plus something small. So it's going up again.

Anyway, so they estimate this-- we're estimating this using the data on how many people listen as a function of N. And they have an IV approach where they use population. You might overesti-- underestimate the effect of N on markets because you'll see a lot of-- or overestimate it because you'd see a lot of stations entering the markets that have a lot of demand.

Like, there are some cities where people like the radio, like Nashville, and some like Boston where they don't, then you might get more firms in Nashville than in Boston because they're meeting the demand. And so the instrument for the number of firms, they use the population as an instrument for the number of firms. And so they're sometimes looking at how does listening compare in cities with a lot of people versus a few people.

Then the second thing they do is just-- they want to estimate this downward-sloping advertising demand curve. They're going to assume that the advertising prices depend on the share of people listening to the radio and market k.

And so it's just the log of price of advertising in market k is equal to some characteristics of the city, plus something-- if you have a downward-sloping demand in the number of listeners-- listener hours that there are plus omega k. So, again, the idea is the more listener hours you have, the lower is the willingness to pay per listener hour for the commercials.

And then the fixed costs. In some sense, this is like Bresnahan and Reiss kind of approach, that you just-- you are trying to rationalize the number of radio stations in every market. And so you rationalize it by assuming that there are these errors. In this case, they do errors drawn from a log-normal distribution instead of from a normal distribution. And they do the-- estimate two parameters-- the mean and variance of the log-normal distribution. But anyway, that's just like Bresnahan and Reiss.

And the data have information about 3,285 stations, 135 markets. So they have individual station characteristics that they're using in some places. These regressions, like the advertising demand regressions, those are just run-- their advertising data is purely at the market level. It's not at the station market level. So there, they're just using market level advertising prices they get from some consulting firm that provides advertising prices. I think it actually comes from the revenues and rate. It comes with radio stations advertising-- reported advertising revenues.

The raw data, which is going to tell you a lot about what the results come out, is here in the paper. So this graph is as the number of stations in a market increases, how does listening go up? And you might say, OK, I definitely see that positive slope-- more radio stations, more listening.

But if you look at this graph, this scale is really running from 0.1 to 0.18, or 0.12 to 0.18. So when you have 15 radio stations, it's looking like you're getting 14% of the people listening to the radio. When you have 45 radio stations, you're getting like 15 or 16.

So we're tripling the number of radio stations in the market, and listening is going from 14% to 16% or something like that. So this is saying that it does seem like radio station entry is mostly business stealing. And, again, we think that this graph, if you drew the OLS, might be overstated how much market expansion there is. So when you're on the IV, it's not surprising that it says, yes, there's an awful lot of business dealing here.

And then-- OK, yeah. This is just showing you where the instrument comes from. That is, it is true that the cities that are small, the cities of a few hundred thousand, tend to have only 10 or 15 radio stations. And as markets get bigger, you're getting more radio stations. And so that's the first stage.

And then the advertising prices. As the listening share increases-- this is per listener hour advertising price or advertising revenue-- advertising revenues per listener hour are clearly declining as the listening share is going up. And the model for this would be there are some things that you need to advertise on the radio. And so if there's very few radio commercial slots available, we pay a lot for them.

And then beyond that, you're just left with playing radio ads for things that you could advertise anywhere else. And the advertiser-- there's just fewer advertisers who really want those ads [INAUDIBLE] the prices after the decline.

So what do they get? So anyway, they estimate sigma close to 1, entries mostly business stealing. Therefore, there's going to be too much entry.

Almost shocking results here on how much excessive entry there is. So welfare under free entry is $5 billion. A social planner could increase social welfare by almost 40% by restricting the number of radio stations.

If we even had a monopolist-- so there was just one firm that had a unique license to operate all radio stations in each market, and that monopolist could operate as many stations as it wanted, it would say we would do almost as well under monopolist as we do under the social planner. So the estimates imply that we want very few radio stations.

A monopolist wouldn't open just one. A monopolist would open three or four or five stations per market. But in some sense, the model is saying that there's just an awful lot of wasteful duplication of fixed costs, so we could get by with many fewer radio stations.

So in their data set, there are 2,509 within metro stations. And they basically estimate that we have four times as many as we want, that optimal would be to only have 649 stations.

What would happen in this world where 3/4 of the radio stations disappeared? Advertising price would go from 277 to 336. So it goes up, but only moderately. Listening share would go from 12.9% of the population to 9.3%.

So it is true that more than a quarter of all radio listening would stop when all those stations disappeared. But the big social welfare gain is the fixed costs of operating all these radio stations drops from $5 billion to $1 billion.

So I just have to say, this does make you wonder a little bit about this. Does this seem like some Marvel superhero movie where there's this evil guy who's like, I'm going to destroy 3/4 of all the people on Earth, and then everyone's going to be better off somehow. So these guys, they're very well [INAUDIBLE] and destroy 3/4 of all the radio stations. And yes, fewer people are going to listen, but we're going to save all this money, and we're all going to be better off here.

One thing they note is that it is true that the advertising listening share goes down from 12.9% to 9.3% Again, remember, this is a calculation where there are only two types of people in the world. There are radio stations, and there are advertisers. And the people are just intermediaries who are just being used to generate ad revenue.

They ask, what happened-- how important would people have to be and people's enjoyment of radio have to be in order to not want to kill 3/4 of the radio stations? And my take on the number is that it's really kind of small.

So the advertiser revenue per listener hour, if I'm reading the data correctly, is about 4.2 cents per listener hour. That is, if you think about it, like you listen to radio, how often you actually buy anything that you hear in a radio ad? How much would someone be willing to pay to have you hear those radio ads? The answer is almost nothing.

So to offset all these losses, you would need consumer surplus to be about 13 and 1/2 cents per listener hour to say we don't want to kill 3/4 of the radio stations in the United States. And again, I-- again, I don't know how much people are willing to pay for listening to the radio.

It is true that satellite radio doesn't get that many customers. And I don't know how much Spotify charges per listener hour. But 13 and 1/2 cents is not such a big number. And so, in some sense, you wouldn't need customers to be all that important and not want to get rid of all the firms in this way. Any questions on that?

Oh, these are just-- yeah, anyway. Had the graphs up there. So yeah. So the monopolist different numbers-- so like, monopolist is nearly as good in their market. The monopolist gets all the way down to just 341 radio stations in 135 cities. So we're going down literally less than three radio stations per city. And again, the listening share dropped substantially under monopoly.

So some comments I wanted to make. In some sense, I didn't need to actually do this paper-- do the empirics of this paper. This is-- as I said, they model this as a homogeneous good. I already gave that theorem that once when you have a homogeneous good, entry is always excessive.

So whatever parameters, sigma and gamma and whatever they've estimated, we know that this model is going to say there's too much entry if the advertisers-- if the people are homogeneous goods being sold to advertisers.

It's definitely true that they're estimating how excessive it is by estimating that sigma parameter. And if that sigma parameter had to not be zero, they would have gotten a much less excessive entry. But the model is always saying that entry is excessive and doesn't have estimates of the things that will let you say there's insufficient entry.

What would those things be? For instance, if you could estimate how differentiated are the goods-- are people who listen to classical music versus top 40 music valued differently by different advertisers? Do the advertisers gain from match quality? That would be one thing leading us to want more stations besides the consumer surplus.

The paper also is also very limited in letting quality vary across stations. I think it is true that there are big stations that have very big fixed costs. There are a lot of college radio stations that have very small fixed costs and very low-powered antennas.

And you imagine that it's not the case that we cut out 3/4 of stations, we cut out 3/4 of the fixed costs. It may be that all those little stations that no one's listening to also are there because they all have very little fixed costs, and we would like to do that.

The other thing that I think they don't have the data to do, but it would be-- I think a big question in these media markets is, what is the number of advertising minutes per hour, and how is that optimal relative to consumer preferences and advertiser preferences?

And it would be very nice to have more on-- know more about how do advertisers choose the number of minutes per-- how do radio stations choose the number of minutes per hour, and what does the trade-off on that look like

And you can imagine that we don't get price variation on the consumer side, but you do get minutes of advertising variation on the consumer side. And so you could imagine that exploiting that minute of advertising variation can be an interesting way to get better at what are the-- what's the consumer surplus of the consumers who are listening to the free radio?

And I should say there are a number of papers that do have advertising data to try to address some of these questions. But I think it's still an area where we could do some interesting work.

So then the final thing I want to do is a couple of papers that are, in some sense, arguing that often, these entry models are doing the wrong model of entry. Most of the models of entry that I discussed in Tirole's textbook are of this variety. We have N for-- we have some large population of firms, and they're all choosing in or out, paying a fixed cost. And then once they enter, there's this symmetric competition between them.

I think we know, though, that-- I've got a radio is an example-- that it's not symmetric competition. There are some big firms, and there are some small firms. And what's in these next two papers is thinking about, even when you have products--

Like in radio, we think there are high-quality, or there's popular music stations that lots of people listen to. And then there are esoteric stations that not a lot of people listen to with low power. Even when you look at products that seem like they should be homogeneous, it seems like symmetry of competition is not a correct answer. And so these papers on entry that are, again, also papers really about competition, are arguing that this model of entry and then symmetric competition seems like it doesn't fit the data very well.

So anyway, here's an example of homogeneous product-- mayonnaise. I've heard that if you do blind taste tests, people have a very hard time telling one brand of mayonnaise apart from another brand of mayonnaise. Shouldn't mayonnaise just be a homogeneous good, where there's a quantity of mayonnaise, market-clearing price in mayonnaise.

Anyway, it turns out empirically that that's actually-- that's not true. Well, it's not true that mayonnaise doesn't all taste the same. It's not-- what's not true is that people treat mayonnaise like a homogeneous product when they're deciding what to buy in the supermarket.

So, in particular, one of the things that's-- so there's this paper by Golder and Tellis that's a few decades old now, noting that there are many, many leading brands that have been leaders for a very, very long time. Like, when did Coca-Cola become the leading cola in the United States? It's 100 and something years ago.

And like many of the products you can think of have been the same-- the same leader has been there for decades and decades and decades. And Eveready batteries or whatever, like the Pampers diapers. These same firms just persist for decade after decade after decade.

The interesting fact that Bronnenberg, Dhar, and Dubé were bringing out in their paper is that not only is it true that in many consumer goods there are market leaders, but those market leaders differ in similar markets. So this is the mayonnaise graph from their paper.

So there are two main mayonnaise firms in the United States, at least in the-- probably in the South. There's Kraft, and there's Unilever. And Kraft produces Kraft mayonnaise. Unilever produces two brands, depending on what city you're in. They produce something called Hellmann's Real mayonnaise. They also produce something called Best Foods mayonnaise.

You can tell from the jars that Hellmann's and Best Foods mayonnaises are-- they're the same mayonnaise produced by Unilever. Just in some cities, they sell it as Hellmann's, and some cities, they sell it as Best Foods mayonnaise.

But then what's striking is this is data on-- over the course of half a year, I think-- on what's the market share by week of Kraft mayonnaise versus Unilever mayonnaise in Denver? And you can see in Denver, 65%, 70% of people buy Kraft mayonnaise, and 15% to 20% of people buy the Unilever mayonnaise.

And then you go to Los Angeles. And in Los Angeles, only 25% of the people are buying Kraft mayonnaise, and 60%, I think it's the Best Foods mayonnaise sold on the West Coast. And so it's striking that like, in some sense, in Denver, Kraft is the clear mayonnaise leader. And in Los Angeles, Unilever is the clear mayonnaise leader.

And so there are these different market leaders, and they've been different leaders. And it seems quite stable that there's a leading mayonnaise that everyone buys. And then there's a second mayonnaise that some people buy.

They give other examples in the paper. If you look at ground coffee, I think of as another fairly undifferentiated product. I don't know if you guys are wealthy enough that you don't drink frozen-ground coffee, and you go to Starbucks or whatever. But you buy the frozen-ground coffee in a can-- I don't know how different the different cans are.

But anyway, Folgers' share ranges from 16% in New York to 59% in Des Moines, Iowa. Maxwell House's share is only 4% in Seattle, but it's 46% in Pittsburgh. And, again, these shares seem to be persistent in the data.

So what Bronnenberg, Dhar, and Dubé do to talk about this is they just investigate this as being a really, really long time span echo of what happened when these products were introduced. And so their idea is that in consumer packaged goods markets, people just-- somehow there's some model where people learn to buy a brand. They buy that brand. And either they buy that brand and they get used to exactly what that brand tastes like, and that's what they think mayonnaise should taste like.

Or they just know that I buy Kraft mayonnaise. My mother bought Kraft mayonnaise. Her mother bought Kraft mayonnaise before her. When you go to mayonnaise, you've seen your mother buy Kraft mayonnaise. You buy Kraft mayonnaise. So anyway, that's the story that they're thinking of, is there's just something about people forming tastes.

It could also be something that's less-- it could also be something that's less taste based. It could also be there are some advantages to dominance. Like, if you're the dominant mayonnaise, you can go to the supermarket and tell Albertsons, I want my mayonnaise at eye level. I want you to put that Unilever mayonnaise down on the bottom shelf.

And that somehow the dominant mayonnaise gets put at eye level. The second mayonnaise gets put on the lower shelf. And then people are just always seeing the eye level mayonnaise, and that's what they buy. And so there's just some-- sometimes this is a story of dominance producing dominance and some kind of market power being produced by large market share.

But what they've done is they go back to-- they go back 100 years. And so all of these products, like the Hellmann's mayonnaise or the Kraft mayonnaise, they find where those products were first introduced in the early 20th century, and sometimes in the 19th century. And then they run-- they do two things.

So this is their main analysis where they have 40 brands which they could find these origins. And they regress the market share that brand i is getting in city m. Forget about the c's. It's a constant. So Hellmann's is better than Kraft.

Plus these coefficients based on your distance to the origin city. Are you the origin city? Are you 50 miles from the origin city? Are you 100 miles, 200 miles, 1,000 miles? The thought being that they don't know when all these mayonnaises got to all these different places. But their idea would be that there's-- if we look at United States.

So anyway, I have a map of the United States. And then there's just like Kraft mayonnaise gets invented here in the year 1895. And then their thought is that it probably just spread like this. And so by 1910, it is here. And by 1930, it was here. And so we're seeing the Kraft spreading, and we're seeing Hellmann's spreading.

And so how far you are from where it started is going to be how many other mayonnaises were there before you. And so if you were near the origin, you were there and got to establish the early lead and now dominate. If you're far from the origin, by the time this mayonnaise gets to California, people were eating several other mayonnaises, and so it's stayed behind. So that's the main analysis because they were able to do this for 40 products.

The second thing they did is there were six products where they were actually able to identify at the city-by-city level, when did Miller beer get to Milwaukee? And when did Miller beer get to Chicago? When did Miller beer get to Indianapolis? When did it get to Philadelphia? When did it get to Miami?

So they actually were able to track when each of the brands entered all the different cities. And then they can regress city-specific market shares on dummies for entry order. Were you there before your rival? Which mayonnaise got to Phoenix first is the regression.

Anyway. So here's what their data sample looks like. So again, this is 40 products. And so for each of these 40 products-- let's go to mayonnaise down at the bottom-- mayonnaise at the very bottom.

Kraft mayonnaise-- oh, I was just about right. Kraft mayonnaise started in Salem, Illinois. And it started Salem, Illinois in 1931. And Unilever mayonnaise, which is Hellmann's on the East Coast-- Hellmann's mayonnaise started in New York City in the year 1905. And so what they're looking at is how do how does the share of Unilever versus Kraft depend on your distance relative to these two points of origin?

OK. So distance to the origin analysis finds that current market shares are 18 percentage points higher in the closest cities than in the most distant cities. So the mean share of these products is like 22%. So this is saying that you get like-- you can think of the 18 difference as being-- you get 31%.

Yeah. Let's see. So you get like 31% in the closest cities, and you get 13% in the farthest away cities. So your market share is just dramatically higher in the cities near where you started-- right where you started than cities farther away where you started.

We get this big effect for your origin city. And then it's declining as we get to 250 miles, 500 miles, 1,000 miles. It then seems to be fairly flat, and then it goes down a lot when you get to the opposite coast of the United States. So it's like East Coast products are selling noticeably less on the West Coast. West Coast products are selling-- or West Coast products are selling notably-- considerably less on the East Coast.

So you may think that with our politics, the coasts are actually similar to each other, and it's the middle of the country that's different. I think this graph is saying, no, it's actually the-- it really is the distance. It's like there are your East Coast mayonnaises and your West Coast mayonnaises. There are East Coast bagels and West Coast bagels. And the market shares are most different on the opposite coast of the country.

And then the second part of the analysis about just is your market share higher or low-- higher in the cities that you reach first versus second relative to your main biggest rival? And these numbers are more moderate. They range across products.

So like the Budweiser versus Miller beer, there's not that much difference between Bud and Miller market shares in the United States. On average, whoever got there first is doing a 1.3 percentage point better than whoever got there second. In the mayonnaise example, who got there first versus second is 6.3 percentage points.

Sure. So it is, I think, fairly strong evidence, especially that first part of there just are these asymmetries. And remember, these are events 100 years ago. And so it does seem like there's very, very long-lasting market leadership advantages.

And if we're thinking about entrance, it's not like you enter, and one year later you're going to be on even footing. It's like you're going to enter, and 100 years later, you're still going to have a fraction of the share that the leading firm had 100 years ago on something like mayonnaise, where it's clearly-- it's not like their mayonnaise has just been better over the years. OK.

So then the final paper I was going to cover today-- Bronnenberg, Dubé, Gentzkow. So this is really digging into where do long-lived advantages come from? They draw two main distinctions for where-- I think two main types of stories.

One is what they will call demand side stories or preference side-- preference-based stories. Their preference-based stories are people develop tastes for products that they consume growing up. Like, I'm just used to this product. I like the taste of this product. Or I've just learned from my parents that this is the right product to buy. I just-- I have preferences about what this product is.

The supply side stories are the ones I talked about about things like shelf space or advertising, where if you're dominant, you get the premier shelf space. Or if you're dominant, it pays you to put on TV commercials. You put on TV commercials that reinforce the people's love for your type-- for your product.

And so what Bronnenberg, Dubé, and Gentzkow want to do is say, does it look like demand side, or does it look like supply side? And they have this sort of very thoughtful idea, which is, how about to investigate this, let's look at people when they move.

And so like many of you have moved to Boston. And the question is, once you've moved to Boston, do you guys buy now the Boston mayonnaise brand? Or do you guys buy the mayonnaise brand that was big where you grew up?

So many of you probably-- I think-- I'd have to guess that Boston is a Hellmann's mayonnaise city, given that, again, Hellmann's started in New York, and Kraft started in Illinois. Again-- and there was a 6 percentage point difference.

But anyway, I would guess that Hellmann's is the Boston mayonnaise. And the question is, do you guys, when you move here-- I don't know. I guess not that many of you are from the Midwest either.

But when you guys moved here, did you start buying Hellmann's mayonnaise and give up your Kraft mayonnaise? Or do you continue to buy the Kraft mayonnaise, which is totally available in Star Market and Whole Foods and whatever. Well, I'm not sure if it's in Whole Foods. But anyway, that's their idea.

And you might think, well, that's a great idea. How would I ever find this out? And this is where, again, you get the recognition of the importance in our profession of being very well-funded and well-connected sometimes.

So what-- University of Chicago has a long-standing relationship with Nielsen's. Firms like Nielsen, besides doing things for people they have relationships with, also do things for people who have money to pay them to do things.

And so what Bronnenberg, Dubé, and Gentzkow got Nielsen to do is-- have I talked about Nielsen Homescan Panel? OK. Nielsen Homescan Panel-- if you want it-- it's actually-- a lot of this data is available for free through the University of Chicago website.

But anyway, Nielsen Homescan Panel is they've got a panel of thousands of people across the United States. And what these thousands of people do is they all have a part-time job. Their part-time job is recording everything that they buy, and they're employed by Nielsen to do that.

And so they're given like a supermarket scanner. So you go to the supermarket. You buy stuff. You scan it all in the scanner when you're there. And then you bring it home, and as you take it out of the bag, you scan it again. And you just scan it in so that Nielsen knows everything that you bought.

And you're supposed to do this everywhere you go. You go to Target, you come back, you scan it. You go to the CVS. You come back. You scan it. You go to the supermarket. You come back. You scan it. You basically just are home scans.

What these people do is they buy stuff, and then they just record everything they got. And then they're also recording where they got it. And so because Nielsen is already in the Star Market, and Nielsen is already in the Whole Foods, you bring it home. You scan it.

They know what you paid for it because they know where you bought it, and they know what that store is charging for it. So you get this sort of data set of what is it that people at the individual level are purchasing. And you can also get aggregate market shares out of this.

Nielsen Homescan Panel did not tell them-- it tells you where the people are located, that this person is in Boston. It doesn't tell you that they grew up in Illinois and moved to Boston when they went to college at age of 19. So what they got Nielsen to do is survey the members of the Homescan Panel and ask them.

So what state were you born in? When did you leave your state of birth? When do you-- how long have you been in your current location? And ask you other questions like, what's your gender, or what's your highest level of education completed, and so on.

Normally, you would think you send a survey like that to people, it's very hard to get them to respond. But since this survey was sponsored by Homescan, who is an employer of all the Homescan people, they got a 65% response rate on their survey asking this information.

They're able to merge that information together with the people's purchasing patterns, and they get this remarkable data set of 80,000 individuals living in 49,000 households. And they know their life history of moving, and then they can examine what those people purchased.

And I think it's an interesting contrast with Bresnahan and Reiss, where we have like-- we're opening up the phone books, and we're counting there are three dentists or whatever, and that's for 200 different cities versus this 80,000 people, you've got basically their shopping lives laid out in front of you, as well as some of their life history.

Nielsen classifies UPC codes by product category and brand. In this paper, they've picked 238 different modules. So a module could be mayonnaise. So they pick 238 modules. They pick modules where they have at least 5,000 households purchasing that item and where the top two brands are owned by different firms.

So, for instance, in the cream cheese module, the number one cream cheese is Philadelphia cream cheese. The number two cream cheese is Philadelphia Light cream cheese. So they would throw out cream cheese. But in mayonnaise, number one mayonnaise is Unilever mayonnaise. Number two mayonnaise is Kraft mayonnaise. They keep mayonnaise.

Anyway, the current market share data is also supplemented with historical data at the state level that's obtained from some published surveys that asked people about, what brand of mayonnaise do you normally buy? And they found these archival things from the '40s up to the '60s. I won't talk so much about that part, though. Actually, they also do advertising data. I'm not going to talk so much about what they do with the advertising data.

But anyway, so here is part of their product list. And, again, it's 238 things long. But so all of these are-- well, let's see. Down here at the bottom, I've got beer. I've talked about beer before.

So the number one beer is Budweiser. The number two beer is Miller. What they do is they record the aggregate purchase share. So aggregate share here means share of product one. Aggregate purchase share is share of Budweiser and the fraction people buying Budweiser or Miller. So 64% of the people who buy Budweiser or Miller buy Miller.

And then, again, these are the cross-state standard deviation. So it's 64 on average. But in different states, it's different. So it's 64 plus or minus 19 for Budweiser versus Miller.

Some of the products, like of bread, Nature's Own versus Sara Lee Soft and Smooth, it's like 50/50 between those two plus or minus 32.

Some of them there's very little variation. So antacids-- Prilosec and Rolaids are the top two antacids. Prilosec is a 71% share in that comparison. It's a more effective drug, I guess. And it's 71 plus or minus 8.

I look at some of these, and they seem like things that are very similar products. Like, I think I would be shocked if Comet abrasive cleaning powder and Ajax abrasive cleaning powder are very different. Some of them, like the bakery products, maybe it seems like they're not as homogeneous as some of the other things.

So what does shopper's homescan-- or these homescan shoppers in the United States look like? I don't know how this compares to the US population as a whole. Someone must know that. About 27% of the households in their sample are migrants.

I guess, obviously, almost everybody I know is a migrant, which I guess tells me that migrants are unusual-- that we're somewhat unusual relative to the US population. But 27% of the households were born in a different state where they now reside. 16% were born in a different census region-- so conditional on having moved states. Many people have moved fairly far.

Obviously, this is-- you'll notice people don't really move to the Northeast. People are born in the Northeast, and then they stay in the Northeast. Very few people, especially from the West, move to the Northeast, given the weather. More people in the South move to the South from the Northeast or from the Midwest.

Ages-- the people in their sample tend to be in their 40s, 50s, 60s. You get very few people in their 20s who want to have a scanner and scan everything they buy at the grocery store.

Age leaving state of birth-- there's an awful lot of variation in that. You get a lot of people who left before they were 10 years old, a lot of people who move in their 20s, not surprisingly. And then you do get people who also move in their 60s when they're retiring.

So what's the main analysis? The main thing is to ask, let's look at these people and ask when they move, how far do they move from their birth place purchasing patterns to their destination purchasing patterns?

And so, they write mu sj for the share of people who never move-- the share of nonmigrants in state s who buy product j. And then mu s prime j is the people in state s prime who never move-- what fraction of them buy product j?

And so what I'm doing is I'm thinking of someone who moves from state s to state s prime and saying, how does their purchasing compare to-- how far in the direction of s prime has their purchasing-- how far along the way from mu s to mu s prime has their purchasing moved, assuming they started at mu sj when they were in-- before they moved 20 years ago when I don't observe them?

So the question is, if this-- if you still purchase-- if these migrants still purchase exactly what they purchased-- well, if they were representative of the population in state s, and still purchase what they purchased before they left, beta would be 0. If they all purchase are exactly representative of the nonmigrants in their new home and purchase what people in the new home purchase, beta would be 1.

So beta is like, what fraction of the way have people moved from looking like where they were born people to where they now live people? And we're going to basically have a regression estimating beta ij as a function of the age that you were when you moved and the time that you've been in your new location.

So you might think if you moved when you were one year old, you've been your location for 60 years, you ought to look like people who grew up in that location where you moved when you were one year old, other than maybe your parents taught you about what mayonnaise to buy.

Whereas, if you're 60 years old, and you moved at age-- if you moved at age 60, and you've been there for six months, you might expect that those people would look more like the origin people, at least on the taste side.

And then they do this via weighted least squares. Because if you think about it like-- if mu s and mu s prime are almost identical-- we're getting almost no information from that market, and it's where mu s and mu s prime are very different that we're getting information, and so they do a weighted least squares estimate weighting the behavior of people who made moves to more different places.

Anyway, so they've got a few figures that they do. They've got this one. Sometimes they're estimating this coefficient separately. Again, you've got 80,000 people. You've got a lot of people who move at age 15 and have been there for 35 years.

So they have separate coefficients for every age at move and every year since move. I find this one hard to see, though, in three dimensions. It's hard to see how high these things are versus the slope.

So easier to read is these ones here. They just go-- they collapse it down to one dimension. And what this is saying is that an awful lot of it seems to be supply side.

Because if you look at people who have just moved-- so it's zero to four years since their move, they're immediately like 60-something percent of the way to the new purchasing patterns. When you look at people who moved 40 or 50 years ago, they are like 80% or 90% of the way to the purchasing patterns.

So it seems like there's a big supply side element that makes you immediately on moving here to Boston, you start buying Hellman's mayonnaise. But there is something that's longer term.

And so this could be that advertising playing on you over 40 years, or this could just be your tastes adapt. You go to friends houses. They serve Hellmann's mayonnaise. You think, wow, that Hellmann's is good. I should buy that myself. So there's adaptation over time, but a lot of it happens on day one.

And then here's the age at move graph. The people who moved when they were children-- obviously, the people who moved when they were children have already had to have lived for quite a while in the new location because they were only surveying adults. But the people who moved at children-- you moved at age one. You do retain something of your birth mayonnaise.

So this is you're getting something about mayonnaise from your parents. Because it can't be that you moved at age one, therefore, you remember the mayonnaise of your babyhood. So there is something that you're getting. It's only 0.9, but then it declines in the age it moves. But it stays that even people who moved in their 60s, they're mostly adjusting to the new mayonnaise.

OK. So what do I have here? What does it say? So anyway, people instantly move 60% of the way. This suggests that it is a lot of supply side stuff. Yeah. People become more similar over time. I said that. Betas are related to age and move. So age 60-plus, you're still getting a big movement toward new things.

Oh, number 4-- data on premove migrants. Again, there's this thing of you guys moved to Boston because you're not the typical people where you grew up. You probably don't drink the same coffee that people where you grew up drink. You probably don't drink the same beer that people where you grew up drink. You probably don't-- maybe you don't have the same mayonnaise like people you grew up.

So the question is, is this in part-- is this like 60% figure because the people-- remember, most of the analysis are only observing where people are today. So the 60% figure is that people living in Boston, maybe it's that like-- it's like, you are different. Therefore, you get your coffee at Starbucks. Whereas, people you grew up with didn't go to Starbucks or something like that. And you moved to Boston because you were like the Starbucks type person or something.

They have evidence, though-- they actually look at the month-by-month data. They don't-- can't track you precisely, but they're ask-- they did their survey in October of 2007. And then-- no, sorry.

They did their survey in September of 2008. And they asked people-- like, they asked people, how long have you been in your current state? And what are the answers? The answers are integers. So it's less than 1, 1 to 2, 2 to 3 whatever.

So all the people who said less than-- I'm living in some state in September 2008. I've been here less than a year. They know that they moved sometime between October of 2007 and September of 2008. They didn't ask for the month of the move.

But what they can then do is say, look at these people who moved here within the past year, and look at what mayonnaise did they eat before they moved. And what they find is that before they moved over this year, it's looking to me like the people were only like 10% of the way to their new destination.

So maybe they were more worldly. Maybe they already had more Boston kind of tastes. But they're only maybe 10% Boston taste people.

And then you look at them as they move through the year. And by the time they moved-- like that 60%, 1 to 4 years isn't 60% 0 to 4 years. That's 60% the within the first 12 months after you get there.

So some of these people have only been here one month and are already buying the new mayonnaise. Some of them have been here six or nine months. But it seems that this 65% thing looks like a causal effect of moving. And it looks like-- or maybe 80% of it's a causal effect of moving. Some part of it is like a preexisting, underlying heterogeneity. But then there's this big, big one.

And, again, it's just-- they're just-- 80,000 is a good-- 80,000 data points is a good thing. And 80,000 data points times 238 products is a good thing. And just-- this is vantage of data scale, as you can do these things that would be very, very hard to do with any small [INAUDIBLE]. Questions on that?

So that is what I got today. So next week, again, is theory Monday, empirics Wednesday. Next week, the theory topic is strategic investment. Again, this is going to be a very Tirole textbook kind of lecture on Monday. If you look at chapter 8 of the Tirole's book. It's the longest chapter in the book. So I will just try to go over the basics of strategic investment.

And strategic investment is about thinking of things firms do to affect entry in markets. Because if entry really is the primary determinant of markups, really should want to affect entry and keep firms out of markets [INAUDIBLE]. See you on Monday.