[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: So bounded rationality. So we start that, you know, there's really, if you go back and look through the IO literature, there's this very, very long history of studying bounded rationality. In some sense, it starts at the same time that Samuelson was having of revolution in going towards rational economics, the entire time there was active debate about whether we should be modeling things rationally or not. And if you look at the early IO with bounded rationality, actually the difference from today is there's an awful lot of IO with bounded rationality on the firm side.

There's a lot of skepticism of this idea of firms as sophisticated, profit-maximizing entities that just always do what's best. And so in the 1940s, there was this debate of should we think of firms as maximizing profits or should we think of firms as trying to get a pretty good profit that's secure and not trying to maximize profits? There was also a lot of literature about this team theory or Simon about firms have to process information that's really complicated, just like firms have to manufacture goods, firms have to deal with information.

They have bounded information processing capabilities. They organize themselves to do this, and there are costs of trying to set the right price or costs of-- just like there are costs of trying to produce goods, and we should take that those costs into account. I'm actually going to talk a little bit more about Smallwood and Conlisk on Wednesday, but Smallwood and Conlisk were very clear arguing, when first rational consumers were coming in, that they felt rational consumers were much less plausible than irrational consumers and we shouldn't go all the way to assuming that consumers are these hyper sophisticated, rational people.

More recent literature has really, over the last decades, has focused on abandon the irrational firms and gone with hyper sophisticated firms, but firms exploiting boundedly rational customers. That's a big literature and there's a thought of, you know, this is why we want to have a consumer financial protection bureau or other agencies, because firms are going to try to exploit irrational consumers and that creates inefficiencies.

These inefficiencies can be large because-- like normal inefficiencies are triangles of price are too high, consumers who have a value a bit above cost don't get to purchase things. If you put in boundedly rational consumers and consumers are buying things they don't want or paying much more than something is worth to them, you can get bigger inefficiencies. And I would say there's three different approaches here that you see taken.

One is the rule of thumb approach, following Smallwood and Conlisk, which just says, rather than modeling consumers as rational, let's just say what we think is a reasonable model for what consumers would actually do in a situation like a consumer walks into a restaurant. Do they really have a utility for every different thing on the menu and have beliefs about how much they're going to like everything they don't have and rationally maximize? Or do people just look at it quickly, pick out something they think they like, and buy it, or whatever.

The explicit bounded rationality. In some sense, search theory is a model of explicit bounded rationality that consumers may be rational but have computational or other informational costs, and they do this combination of doing the best thing they can given their computational costs. Search is one example, but there are many others. And then the third branch is the psychology and economics branch where this is building off the kinds of things you would see in a behavioral economics class and modeling consumers as having behavioral biases and thinking about how firms exploit those.

You know, this has been the most active branch in recent years, and so today's lecture, I'm pretty much-- unless I get to the paper at the end I think I won't have time for, today is pretty much going to be about these psychology and economics papers. What goes on. If you take a behavioral economics class? You know, so over the last 25 years, there's been this explosion of work in behavioral economics.

And the way behavioral economics classes normally go is you talk about and read psychological studies that talk about how consumers do things that are not rational, and then you discuss models that people have written to model behavior that's kind of close to rational but not quite rational in a way that would explain the psychological bias. I mean, I guess Drew does a little bit of this in 123.

But it's you take a model that-- take a standard model and then you make some small tweak to it that would allow consumers to behave in some different way, and then get a utility-like model that may have one or two extra free parameters that could display this is nearly rational but with these two extra parameters describing their behavior instead of behaving fully rationally. And there are many, many ways in the behavioral economics literature in which consumers are shown to behave not rationally, and there are rational-like models for most of them.

So why don't I-- and, you know, the idea of how this behavioral economics literature could have an influence is you write down a nearly rational model, and then any model that you could solve with rational consumers, you can solve with these nearly rational consumers because you have a well-specified thing for their behavior which isn't just maximizing utility function. So, for instance, loss aversion. In financial problems, what do consumers do? They're taking some action a.

And in the rational model, if they take an action a, they'll have some wealth w of a which is a random variable, and then they have a utility that's a function of that wealth, and then they maximize the expected utility of the wealth after taking the action. In loss averse problems, what is assumed is that consumers have a reference point, w0, which could be their current wealth, and they try to maximize some expected function of w of a minus w of 0.

So you have utility not for wealth but for changes in wealth, and zero is a, somehow, point that's special. And the simplest version of this is that you have a utility function that looks like this where you care much more about losing money than you care about making money, and somehow losing money is painful to you. Because if you ask-- there are all these experiments, if you ask students which-- would you a 50/50 coin flip, win $101, lose $100?

And almost everyone says, no, I don't want to do the 50/50 coin flip, win 101, lose 100, or something like that. It's not enough of a gain to make it worth losing $100. Matt Rabin has this very beautiful paper on calibrating preferences that says it's just completely inconceivable for a rational consumer to say no to a 50/50, 101, 100 gamble, because if you're going to turn down a 50/50, 101, 100 gamble, you would also turn down a lose 5,000, win a trillion gamble. Because in some sense, to turn down that gamble, you've got to have a very concave utility function and it's got to just keep getting concave.

And if you don't like small gambles, you can't-- obviously you all take million dollar gambles with your lives, like when you choose your careers or whatever. Somehow people treat big stakes gambles very different from small stakes gambles. This loss aversion is one way to explain why people might do that. People often-- I guess, more sophisticated models do something more like this as a utility function. But the whole point is that it does have a kink here where the slope to the left and the slope to the right differ.

When you have multiple goods, you have to assume, like, how do we deal with like, OK, this is the utility of wealth. How would we deal with multiple goods? And there are these-- again, these classic experiments, like you bring people into a lab and you have them play some game and you say, as a reward for doing this, you get an MIT Department of Economics water bottle.

And then you let them have the MIT Economics water bottle for a few minutes and then say, oh, you know, actually, I also have these-- I also have these Lindt chocolate bars. You can have the Lindt chocolate bar instead of the water bottle, you know, if you want to trade or whatever. And then everyone says, no, no, I'll keep the water bottle. And then you run the experiment the opposite way. You give them the Lindt chocolate bar and say, oh, we also have this MIT water bottle you could have or whatever.

Nobody trades a chocolate bar for the water bottle. And so the idea there is that how do we explain that people just get attached to things that they have. The loss aversion, what you've done is you have two different functions like this. I have my water bottle utility function and my chocolate bar utility function. And once you've given me the chocolate bar and I've had it for a few minutes, then it feels like a loss to give up the chocolate bar and it would be a win to get a water bottle.

But you know, a loss of a chocolate bar and a win to a water bottle, that you wouldn't do, and the same vice versa once you've changed the endowments. And I think this is-- people do think that this is how you think about-- in some sense, this is one explanation for how stores try to lure people in with and then have high last minute charges or hidden charges is once you've gotten used to the fact that I'm going to own this TV, it's then hard to walk away from the TV because it's now a loss of one in the TV dimension because for the half an hour you thought you had that nice TV or something. OK.

Fairness preferences. So obviously there's nothing in utility function that says that utility is only based on your own wealth, not on everybody else's wealth in the entire world and people can be altruistic. But I think when we're doing normal-- normally when we're doing most topics, we do stick with preferences that are non-altruistic and just everyone's preferences are based on their own consumption. There's a number of different models out there of preferences that do have some other regarding nature to them.

So for instance, Fehr-Schmidt is the most popular model of fairness in the behavioral literature. What consumers do in a Fehr-Schmidt model is they are trying to maximize their own payoff. So I'm going to write pi here for the payoff rather than the utility. They're trying to maximize their own payoff, but they dislike when other people are getting more than them in a game. And so we subtract alpha times the fact that the extent to which other people's payoffs exceed your payoff.

Fehr-Schmidt also assume that consumers dislike when other people earn less than them. They are other-regarding, and so they put weight beta on getting more money than other people in the game are getting. But Fehr-Schmidt assume that most people have a sense of fairness, but the fairness is more acute when they're the ones being treated unfairly, and so alpha is bigger than beta. Then you put more weight on me getting the small share than on other people getting a small share, although I don't like that either.

And then again, you start with any utility model, and then you could do the Fehr-Schmidt model with two extra parameters alpha and beta, and then you can think about what happens in that model as a function of alpha and beta. So in some sense it becomes just a plug and play thing. Take a standard model, put in the alpha beta preferences, and then you can resolve the model and see what happens as a function of alpha and beta. And you can calibrate alpha and beta to what's in the data.

Most of the examples I picked out to do today in IO, I'm doing what are called quasi-hyperbolic preferences. So in preferences with quasi-hyperbolic discounting, consumers are maximizing, at each point in time, the utility of the consumption they get today, plus beta times a discounted sum of future preferences. So everyone always cares about today more than they care about all future periods, and so your trade-off between should I clean up my apartment 11 days from now or 12 days from now, you're completely indifferent between whether you spend a day cleaning up 11 days from now or 12 days from now.

But should I do it today versus tomorrow? You really want to put it off to tomorrow? So there's some much bigger emphasis discounting between today and all future times than there is between future times that come later. So the discounting between any two future periods is delta, but between the current period and-- you put weight one on today, beta delta on tomorrow, beta delta squared on the next day so that the ratio is beta delta here and only delta here.

Again, so people have preferences. Like I would like to go on a diet, but I would like to go on a diet starting tomorrow. Or I would like to save more money, but I'd like to save more money after I buy this one TV I really like and then I'm going to start saving my money, or something like that. And then there are variants of this model. Sophisticated consumers recognize-- sophisticated hyperbolic consumers recognize that they have these preferences and they want to do things to commit themselves to not behave in a hyperbolic way, or they know that they need to commit their future selves to things so their future selves will actually do what they want their future selves to do.

So it's like, I am going to put something on my phone that shuts me off of social media for some number of hours to commit that future lazy self of mine to actually get work done and not spend time scrolling through social media. The other ones are-- other models, the naive model. The naive model, you're hyperbolic and you don't recognize that your future selves are hyperbolic and you just think, I have willpower. These are my preferences.

It's perfectly reasonable for me to say I'm going to sit here and scroll through social media, and then tomorrow my future self will not scroll. I'll just make the plan that this week I get six hours of social media. I'll use it all up this afternoon, and then my future self will not do any social media tomorrow. And that's my plan for this week. And then the partially naive model, an extension to go from a two parameter to a three parameter hyperbolic extension is I am a beta delta consumer.

I recognize that my future self is also hyperbolic, but I think my future self is like a beta hat delta consumer with beta hat greater than beta. So I just underestimate. I recognize that I'm hyperbolic, but I underestimate the degree to which my future self is hyper hyperbolic. And again, if you think about the delta as already being a parameter we already have, again, maybe just beta and beta hat are two parameters you add to the standard rational model.

Other common examples. I won't talk about IO examples today with them. Imperfect Bayesian updating. Eyster and Rabin proposed this model. There's a lot of evidence that people are not so good at Bayesian updating. They propose this model of imperfect Bayesian updating that's just the simplest one in the world. Your posterior on theta after observing some signal x is some constant times your prior plus 1 minus chi times what the correct preference is, correct Bayesian updating is.

So people just under-update when they get new information. They tend to be overly confident, stick with their priors. Overconfidence bias. Overconfidence bias is people think they are better than they are at many, many things. A classic example is ask a class how many miles is it from Boston to Buenos Aires and write down your best guess for Boston to Buenos Aires. And now write down a 50% confidence interval or a 90% confidence interval.

You admit you don't know how far it is to Buenos Aires, but write down an interval you think there's a 50% chance that the distance from Boston to Buenos Aires is within the interval you write down. And overwhelmingly, I don't know, is it, like, 50% of the people-- like if you ask for 90% confidence interval, 50% of the correct answers are within the 90%-- 50% of the time, the 90% confidence interval contains the answer. 30% of time, the 50% confidence interval contains the correct answer.

It's just people think they are better at many things than they are. Self-serving bias. This is the idea that we all like to think-- this is that we all think we're good at things. Self-serving bias is we all like to think we're good at things. So you make a decision, and then when you're called on to make a second decision, you make a second decision in a way that makes your first decision correct. That is so instead of like admitting-- you make some plan. You're going to go somewhere.

And then you get there and realize, wow, I spent half an hour on the train to get to this place. This restaurant looks way more expensive than I wanted or I don't want to be here. Food looks bad. But once you're there, you go ahead and eat there anyway, because if you don't eat there, you're admitting that I made a bad decision in deciding to pick out this restaurant and come here for half an hour. And then you want to make my first decision correct. The only way my first decision can be correct is if I then carry out my second decision to go ahead and eat there.

Then maybe you don't like the food and you're like, oh, wow that was unknowable, so I just had a bad draw or whatever. But it wasn't an ex-ante bad decision. But the idea of self-serving is we always-- we follow these paths in life where we just keep taking decisions, and each step we take a next step that's consistent with the earlier steps being right, rather than admitting your first step was wrong and I should have just admitted my first step was wrong and said, yeah, I made a bad decision there and gone back and start over or something.

OK, anyway, so many, many things. And in some ways, I think this is a-- field has expanded enough in the last decade that it's no longer as true, but in some ways, once you have this view of-- and this is sometimes why behavioral economics does so well in every field of economics, because if you think about it, before behavioral came into fashion, what had Tirole done in his textbook?

He worked out with rational consumers what happens with monopoly pricing and price discrimination and selling durable goods and horizontal differentiation and vertical differentiation and search costs and so on. You basically take the entire syllabus of this course and it's all rational consumers and firms dealing with this situation. And then if you think about from the behavioral perspective, well, rationality is just one incorrect model of consumers, and there are actually this many or more other models of consumers that would be better models of consumers.

In some sense, what we need to do is take the entire field of IO as it existed in the year 2005 and multiply it out into a matrix. What happens in every one of these models with quasi-hyperbolic consumers, with loss averse consumers, with consumers who care about fairness, with overconfident consumers, and so on? It seems like, how can I think of a paper to write? But in some sense, you can take every paper that exists in the literature and write 14 new papers about it, and so it just opens up this massive field.

And again, is-- what do I have here? Yeah. Are all of those good papers to write? You know, probably not. It may be that maybe you try to do horizontal differentiation, you put in quasi-hyperbolic consumers and then you realize, wait, the hoteling model is a static model. It's all in one period. Therefore, hyperbolic has nothing to do with it.

Certainly loss averse consumers in a horizontal differentiation model matter and consumers could care about fairness and consumers could be overconfident in their thoughts about their own preferences or whatever. But in some ways, there are many, many of these boxes that have something there, and in some ways, I think what you want is there's this opportunity to explore them and think about which of these preferences might be relevant to which of these models we looked at and think about what is there interesting to say.

And I think what a lot of what people are interested in is if you can tell stories where I make some small change to the existing model and now I get a big change in the outcomes, so I don't need an enormous irrationality. In some sense, the original model we started with was very sensitive to rationality. When you put in a small change, it makes a big difference. And you could certainly imagine many of these things, like, it could make a big difference in durable goods models if we had quasi-hyperbolic preferences.

Take an example. I guess you probably didn't get there, but when I talked about my add-on pricing paper about do firms want to post an unadvertised-- post an advertised price that doesn't include some add-on and then hold up consumers for the full monopoly price of the add-on after consumers show up at the firm. Back in the back half of that paper, there's a section that says in a game where the two firms are simultaneously deciding whether to advertise the base good only or the base good and an add-on, which would they do?

And the answer is the firms don't want-- it's not individually rational in that model for the firms to do add-on pricing. However, if you add a small number of behavioral consumers who don't realize how expensive the add-on is going to be until they get to the point of sale, then it can make sense. And so it's just an example where you have a small change to the model, like a small fraction of the irrational consumers can have a big change on the model.

So I'm going to start with just, you know, very early paper bringing out some basic ideas. So this is Della Vigna and Malmendier thinking about how a monopolist would design products or sales mechanisms to exploit hyperbolic consumers. And this is a multi-stage game motivated by health clubs where the idea is that everybody wishes they went to the gym more than they actually do. And so here's their model. So firms offer consumers a two part tariff.

You pay l, which is a sign-up fee to join the gym. And then if you visit x times, you pay p times x. So every time you come to the gym, you pay p to go there. Consumers have these beta delta hyperbolic preferences and they say yes or no, I will pay my upfront fee l to join the gym. And then in their model there's uncertainty. You know you should join the gym and that it's good for you and you don't remember how unpleasant it is to go to the gym.

Like, how long does it take to actually get there, how unpleasant is the exercising there, how gross is the showers at the gym, and whatever, how much are you going to dislike exercising once you get there. And so after you go to the gym, you learn your disutility d that you incur every time you go to the gym. And then you then choose how many times am I going to go to the gym this month. To simplify the algebra, they do either zero or one, but obviously we're thinking of one as being go to the gym every day and zero as never go to the gym or something like that.

And then you incur disutility d plus p times x when you go to the gym. And then you get a delayed health benefit x times b. So you do feel better after you've been to the gym. You're happy that you went yesterday, and that's x times b. OK? And so the hyperbolic preferences come in here where the consumer at t equals one who's hyperbolic wants his future selves to go to the gym because he treats the disutility and the benefit as both having roughly equal weight because they're just both future things, and he thinks the future benefit is worth the future disutility, probably.

But then the consumer who's actually here making the do I go to the gym today or not decision, that consumer is treating the disutility and the money paid as a present thing and the health benefit as a future thing, so that lazy future self is tempted to not go today and say, OK, the tomorrow self will be the one to go to the gym, not the today self.

So what are some of the observations? First monopoly observation is that in this model, we get distorted pricing where p is distorted down. In their paper, they have a costs, marginal costs, c times x of having people go to the gym. I mean, c may be small, but the gym does-- they have to wash your towel. They may wipe down the equipment. They have some capacity levels, so they incur some small costs c when you come. But so the main distortion we get is we get p less than c.

This can be going to the gym is free. You don't pay when you show up on your gym membership. Or it could be that p is even negative, that you don't pay and they give you a free juice or they give you some other things for free that they don't charge you for when you're at the gym. Why? Why do they distort p down? There are two reasons. First is let's think about the sophisticated hyperbolic model where this consumer knows that his future self is going to be lazy.

So that means he knows there's going to be a problem of, I sign up for the gym membership. This guy is not going to go, and therefore, I'm not going to get the health benefit. So what happens is the sophisticated consumer wants to commit his future self to actually go to the gym. And how does he commit his future self to go to the gym?

He makes the price. If you offer a contract where the p is negative-- there's all kinds of free stuff offered at the gym and they don't charge you any money-- then you're thinking, that lazy future self of mine is going to go to the gym because he's going to want to get the negative $10 fee for going to the gym, and then that's going to make him willing to incur the disutility to get the benefit. And so the sophisticated consumers want to have a contract-- and they're willing to pay a lot of money up front to get the contract with the low p that's going to force their future self to go to the gym.

So that's the sophisticated argument. If consumers are naive hyperbolics instead of sophisticated hyperbolics, you also want them-- you also want to do p less than c. And the idea here is that the naive hyperbolic consumers misestimate how often they go to the gym. They say, OK, you know, I know b is bigger than d, therefore, if b is bigger-- if I know b is bigger than d and the price is zero, my future self is going to go to the gym every single day.

And then what you can say is, so if you have this person who thinks they're going to go every-- so suppose that we have the benefit b is bigger than the disutility d, but the benefit b is less than beta times delta. OK? So this is the case where the future self thinks they're going to go to the-- Sorry. This is d, not delta. The future self thinks they're going to go to the gym, but they're not actually going to go to the gym.

Because you think that they're going to be rational and trade off benefit versus disutility, but they're actually going to trade off disutility versus beta times benefit. So this is-- oops. So beta times b is less than the disutility. OK? Well, so here's what you can do. You can tell that consumer, OK, look. What we're going to do is go to the gym and instead of charging you cost every time you go, I'm going to give you $10 less than cost every time you go.

And so you're going to save-- you're going to earn $300 in these bonus payments for going to the gym every day, and so for that, for the low, low price of $299 a month, I'm going to give you that $10 every day. And the naive hyperbolic consumers are like, wow, that's an awesome deal. I have a negative gym membership, so they only pay 299 up front and then I get $300 back. But then obviously, you know, you're never going to actually-- if the negative 10 is not enough to get you to go, you're never actually going to get to go. You've paid 299 up front for something that was worth zero to you. And because you're going to overestimate your gym usage, you then end up-- you want to charge them with that kind of scheme.

OK. Third observation here is in this sophisticated model, the p is socially efficient. There isn't actually distortion. You're charging them less than cost, but the price that you charge them, even though p is less-- we normally think of p equals c as what gives us full social efficiency. Here it's actually p less than c that gives us full social efficiency. Do I have that on here? Yeah. Why is that?

You know, in some ways this is like back to lecture one of this class. If you remember in lecture one of this class, I had if you have a firm choosing product quality, they choose the optimal quality for the marginal consumer. In the Della Vigna-Malmendier paper, all the consumers are ex-ante identical, so the marginal consumer is the average consumer. Therefore, what you do is you always choose the optimal product quality and just extract all of that with a fixed fee. That's what this model is here. It's like p is just a dimension of quality.

A lower p gym contract, a lower per visit contract, is a contract that has a higher quality, and therefore, you just choose the optimal quality for the consumers and then you extract all the fixed fee. It's fully efficient from the point of view of the sophisticated hyperbolic consumer who is the one buying it. And more generally, that theorem from the first class was that it's efficient in terms of maximizing the sum of-- it's socially efficient, in terms of maximizing the sum of profits of the firm and the utility of the consumer.

But it wasn't really the utility of the consumer that mattered. It was the willingness to pay of the consumer. And so if you make the willingness to pay different from the utility, then a monopolist serving a homogeneous population is going to choose the quality that maximizes the sum of profits and willingness to pay. And so it's going to be efficient in that sense of willingness to pay as utility. And so that's what we have in the sophisticated hyperbolic consumer. You're always going to maximize the utility that you get.

And this is going to-- in some ways, a general lesson here is that it's not like competition would force you to eliminate these distorting consumers. Competition forces you to exploit behavioral consumers because competition forces you to design a product quality or a contract that maximizes willingness to pay and then give all the rents back to the consumers. So whether it's monopoly or perfect competition, firms will always exploit behavioral consumers because they're just always trying to maximize willingness to pay.

And it doesn't matter that willingness to pay is utility. You're just trying to maximize willingness to pay. If willingness to pay and utility are different, you maximize willingness to pay, which is exploiting people. So they have one of the first behavioral empirical papers trying to analyze this idea in a data set on health club memberships in the Boston area, and the idea is it's basically a reduced-- one of these reduced form papers where we try to say, what are predictions of the rational model?

What would be not necessarily predictions of the hyperbolic model, but directions in which we think pricing of hyperbolic consumers would differ from pricing with rational consumers? And what they're trying to do is reject the rational model of health club consumers and show that some kind of hyperbolic-- it's departing from that hyperbolic direction. So anyway, the data set has 7,000 members at three Boston area health clubs.

A lot of what they have is just a data set on the usage of these health clubs, and then they do also have a survey data where they stood outside the health clubs and just asked people survey questions about how much they use the gym and stuff. So first, inconsistent with the fully rational model. Customers who sign up for the gym just pay too much. One of their plans their health club offered was a $70 a month unlimited gym pass. You just pay $70.

You can go to the gym as many times as you like. Those who sign up for the $70 a month plan averaged visiting 4.3 times a month, and so they were paying, like, $17 every time they went to the gym. The health club also had a plan where you could just get a coupon book, which was 10 coupons for $100 where you just pay $10 on a per visit basis. Clearly it's not fully rational model. You should not be paying $70 a month for a gym you're going to go to four times. You should just buy the book that gives you the 10 visits and just use those coupons.

So either the consumers are misestimating how often they're going or the consumers value the commitment and are buying the gym membership because if they didn't buy the gym membership, they would go zero times instead of four times if they were actually going to pay $10 every time they went. So we don't know which of those, but both of those naive and sophisticated hyperbolic models predict that consumers should be in this situation where they're paying too much ex-post for the gym visits.

Second thing of showing that consumers are not fully rational is this company had two types of plans. Besides the daily, the per visit pass, they had a monthly pass and they had an annual pass. The monthly pass might have been like, you know-- I forget the numbers. $70 a month. But then you also may have had a choice of you could also pay upfront $700 and get a full year membership, paying less per month if you're gung ho enough on the gym that you would pay $700 fee.

And they note that people who pay the $700 fee up front do go to the gym more than the people who pay the $70 a month. So it's as if the people who buy the annual pass know that they are people who go to the gym and they're willing to pay the $700 check up front to get the annual membership. But then one of the interesting observations is that the monthly plan, they charge to your credit card every month until you do something to get them to stop.

The $700 plan, you actually have to-- it didn't charge $700 to your credit card a year from now. You had to actually go in and sign up for another annual membership. And what they find is if you look 13 months later, people who signed up for the $70 plan are more likely to be in the gym than the people who signed up for the $700 plan. And it's because these people just got auto renewed if they didn't do anything, whereas these people had to go in and actually decide, yes, affirmatively, I do want to still be signed up for the gym.

And so it's sometimes telling that sometimes these people aren't paying attention. Another statistic about this was that on average, their consumers spend $187 after the very last time they ever went to the gym. So if you look at people who quit the gym, they've paid $187 since the very last time that they visited. This is saying that people aren't paying attention-- aren't fully paying attention and the gym is designing a payment scheme with the auto renewal on your credit card to exploit that.

OK. Other observations that suggest the partially naive model may be better than the fully rational explanation for this. You ask people in surveys, how often will you go to the gym? And people say that they're going to go to the gym more than twice as much as they actually go. So you ask the people who have-- you ask people who have signed up for the plans, how many times do you think you're going to go? And they get that wrong.

And then also, there's a correlation between lower usage of the gym and taking a longer time to cancel after you stop going to the gym. And so this is a sense that there seem to be some people who are more rational and some people who are less rational, and then the less rational people do multiple things that seem irrational together, like the not canceling and the not going so much. Again, they don't estimate a specific behavioral model and show that specific behavioral model can rationalize all the things, but it's saying that the rational model just doesn't work. OK.

OK. So I thought next thing I would do is cover in more detail one theoretical paper to just give you more of a sense of how these papers work. So this is a paper by Heidhues and Koszegi about credit markets. And I think there's a feeling that many people have too much debt. In the last financial crisis, there was a view that people have too much debt. There's also an observation that many firms offer credit card deals that look like things like sign up for our credit card, pay 0.9 interest for the first six months to pay off your other debts, and then it goes up to 19.8%

And there's this use of these teaser rates, and then people think that, are those teaser rates designed to get people to take on more debt than they should take on or more debt than they want to take on? And so that's the idea of the Heidhues-Koszegi paper is to think about, would quasi-hyperbolic consumers-- would the market exploit quasi-hyperbolic consumers in a way that creates them taking on too much debt and would it explain why it is we see that these contracts that have initial upfront teaser rates followed by higher rates later on?

So in their model-- and this becomes-- it's a nice paper in that it's a paper that shows there's this very, very sharp discontinuity where, with fully sophisticated quasi-hyperbolic consumers, one thing happens and then you make even an epsilon difference between their self-recognition and full self-recognition, you get a big discrete jump in behavior with an epsilon change in the behavioral assumption. So anyway, we have this agent, they have the beta, beta hat one preferences.

So just we're eliminating discounting because t equals three. So the consumer is going to borrow some money at t equals zero. Think of this as taking out a car loan. I'm borrowing some money c to buy a car, and then I buy a durable good and I get utility from the durable good at t equals one. So there's some sense there's a time zero where I pay the money and then a time one where I enjoy the durable good, and then time two, I could still be enjoying the durable good. They do that the durable good is already worn out and I still just have my debt for my car after I've already sold my car.

But anyway, there's I borrowed money. I get utility of t equals one from my durable good. At t equals one I'm making my car payments, and then at t equals two, I'm still making an additional car payment. So they have this conflict of interest between how the different selves will view the car payments versus the enjoying the car. So from the point of view of the time zero self, which is before he's going to enjoy the car and before he's going to make the repayments, time zero self's preferences are c minus k of q minus k of r.

So there's this convex function about your repayment. You have this convex disutility of a repayment because the more you have to repay in your monthly car payment, the more you're cutting back on other things, and the more things you cut back on increasing disutility for cutting back on food as you're cutting back more and more. And so the time zero self utility is c is the enjoyment of the car minus-- all these things would be multiplied by beta because they're all in the future, but we ignore that fact because they're all in the future.

So from the time zero self, these things are all comparable, and it's c minus k of q minus k of r. From the time one's perspective self, time one cares about the car and the repayment made at t equals one. By the time one self discounts the pain he's going to suffer repaying at t equals two because that's a future thing. And then by the time it gets to the time two self, the time two self dislikes the repayment that the time two self is being forced to make.

We're going to assume that the credit market is perfectly competitive. It has zero interest rates. And firms can offer these option contracts where you pay, you get some amount of c is given to you up front, and then you have these repayment plans you can choose among, q sub s, r sub s, and at t equals one you choose how you're how you want to repay. Do you want to make the immediate repayment at t equals one? Do you want to delay your payment and spread it out over t equals one and t equals two? Or maybe they have many, many different possible plans. OK.

Suppose we have a rational consumer. What does a rational consumer do? Well, the rational consumer, you're just maximizing this, and there's going to be a break even constraint in the competitive market that is going to have to be c equals q plus r. The competitive interest lenders are going to always offer you the full amount of the repayments as your upfront price of the loan. So we're just trying to maximize this.

So we're trying to maximize c minus k of q minus k of r subject to the zero profit constraint. OK. What's the solution to that? Lagrangians are all ones. You differentiate with respect to q, and the repayments are going to have to satisfy the marginal disutility of repaying a dollar when you have a payment of q is equal to one and the marginal value of repaying a dollar when you're paying r is equal to one.

And so you repay evenly, setting these repayment rates so that the disutility of paying that dollar is one, and then you set c equal to q plus r. Anyway, the k prime at q equals one pins down how big the loan should be, and then you just repay evenly over the two periods. And the rational loan contract doesn't have to be some complicated thing. The rational loan contract just says, you know, here you go. You get no choices. You get c first best, and then you make those payments q first best, r first best. There's exactly one repayment option. That's what you have to do.

OK. Imagine that the consumers are sophisticated. So you have sophisticated hyperbolic. So the consumers are hyperbolic, but they're sophisticated so that they have beta delta preferences, but they recognize their future self as beta delta preference. Then the time zero consumer, though hyperbolic, the time zero consumer here is not getting the benefit. It's the time one and two consumers who get the benefit of the car. So the time zero consumer has essentially rational preferences over what he wants the future selves to do.

So in a sophisticated hyperbolic model, a sophisticated hyperbolic model, the monopolist is still going to offer the rational contract. And so you're going to commit your future self to make those two even payments, because that's what is best for the future selves to do from the point of view of the time zero purchaser. So in some sense, a lending market sells commitment to the time zero consumers who want commitment and lets them overcome their hyperbolic preferences.

Again, they get their first best. We get the first best buy just committed to the contract is the first best. That's the only option. Then the interesting theorem is what happens when beta hat is different from beta. The result is that when beta hat is greater than beta, you think you're even epsilon less hyperbolic than you are. There's a discontinuity in the contract. And instead of offering just this straightforward no option contract, firms now offer an option contract.

The contract looks like you get some amount c up front, and then you have two repayment options. One is to repay everything at t equals one and get a very good deal repaying everything at t equals one, and then the other one is to spread your repayments out over t equals one and t equals two, and make an even larger repayment at t equals two than you make at t equals one. So you can think of it as like you have these-- you know, they give you money up front.

Let's suppose they give you $30,000 up front and then they have these two options for repayments. One is to repay 25,000 at t equals one and zero at t equals two, and the other is to repay 15,000 at t equals one and 25,000 or 20,000 at t equals two. OK? So the first best in my simple example would be they loan you 30,000 and you repay 15,000 in each of the two future periods. That's what the first efficient thing would be in this model.

But instead what they do is they actually loan you-- actually, they do following, they loan you 35,000 and they give you two options. Either repay 25,000 at t equals one and nothing at t equals two, or repay 15,000 at t equals one and 25,000 at t equals two. And so the contract has two different distortions. One distortion is that the consumers end up taking on more debt than they want-- take on more debt than they want, more debt than they should in a rational contract and get a bigger loan and buy a more expensive car.

And then the second distortion is that they end up choosing the wrong repayment option, where you give them this option, they think they're going to choose this. Wow, that's a great deal. If I just pay everything back next period, I'm getting an awesome deal repaying my car loan. But then in fact, once it gets the next period and trying to come up with 25,000 all in one year is too painful, you end up spreading it out and saying, OK, I'll pay 15,000 today and then 25,000 tomorrow.

So we get these two different distortions, the distortion of loans that are too big and loan payments that are distorted with a balloon payment at the end that's larger than the payment in the beginning, whereas spreading the payments evenly is what would minimize the customer's disutility. OK. And, you know, it's a model with perfect competition. So with perfect competition, the lenders-- actually, so I should have done this with 20. In this model, the lenders, zero profits.

You get zero profits for the lenders, so the amount that they borrow is exactly equal to the repayment. So it's not like lenders make any extra-- lenders don't make any money. It's just utility is being destroyed by having consumers get a larger loan than is rational for them to get and pay the money back unevenly in a way that makes them suffer more disutility from the unequal payments. So questions on the finding of what they're trying to show?

OK. So what I thought I would do is I would actually just go through the-- it's a standard type of mechanism design argument. It's pretty simple, so I thought I would do that. So first way to think about analyzing a model like this is we can restrict ourselves to contracts that have just two options, because all that really matters is what does the consumer think they're going to repay when they sign the contract and what do they actually repay.

So there's one option that needs to be there, which is what they think they're going to pay. There's another one, what are they actually going to pay? You can put 10 different options in there. The other eight don't matter. OK, so we can think of a general class of contracts as being how much money does the lenders give them up front, what's the repayment option that they think they're going to take when they sign the contract, and then what's the repayment option that they're actually going to take. OK.

The lenders are competing to make the most attractive offers, which means it must be that the offer that gets made and accepted in equilibrium solves this problem. We're maximizing the utility that consumers think you're giving them subject to three constraints. The first is the individual rationality constraint of the lenders, which is the lenders have to earn zero profits, so the amount that they give the consumer up front can't be any bigger than the repayment that they're going to get in equilibrium, which is q and r.

There's the IC constraint for the time zero self, the one making the purchase. The time zero self has to think that they're going to repay q hat r hat. And so it must be that the disutility that you get from paying q hat r hat, which is k of q hat plus beta hat k of r hat, that's how you think your future self is going to be hyperbolic, that has to be less than k of q plus beta hat k of r. So you think, wow, this one is such a bad-- IC zero says this is such a bad deal relative to this one that my future self will pay this one.

And then it must also be that when we get to t equals one, from the t equals one self, when the t equals one self is comparing these, the t equals one self says, wow, this is painful. I'm going to choose this one instead. So those are the two IC constraints. The t equals zero self has to say, I think I'm going to take the q hat r hat option, and then the t equals one self has to say, that q hat r hat option is painful. I'm going to take the qr option instead. OK? So we have this fairly simple problem that initially has five unknowns, c, q, r, q hat, r hat.

And as you often do in mechanism design problems, the way to solve these things is to start by eliminating as many of the constraints as you can or turning the inequality constraints into equality constraints and eliminating variables. The first variable I'm going to note is the IR constraint is always going to be binding. You're always going to set c equals q plus r. You're going to set c equals q plus r, because giving the consumer more money up front makes the consumer happier, and this is the only constraint where c appears. So clearly we can really drop that out and just say whatever q and r we choose, c is going to be q plus r.

OK. Second observation is that this one also has to hold with equality. If this constraint was non-binding, I can make q and r bigger. And if I make q and r bigger this constraint, if this was non-binding, I'd make q and r bigger, this would still be satisfied. This would be satisfied, this is satisfied, this doesn't involve q and r. So I'm always going to want to make q and r as big as I can until this constraint gets satisfied. So what limits these repayments is that the consumer at t equals one has to actually be willing to make them. So anyway, so this one holds with equality.

So once you know that this one has an equal sign in it, think about what does this constraint equivalent to when this one holds with equality? Well, if this one holds with equality, I can just subtract the two left side-- the left side here from this, subtract the right side here from this, and both things are still true. And so the top inequality still holds, and these things cancel, and these things cancel. And so what I find is that it must be true that beta hat minus beta k of r, the right side, this minus this, is bigger than beta hat, minus beta times k of r hat.

So the IC constraint. So once this holds with equality, this constraint is simply equivalent to r hat being less than r. OK? So IC zero is just this holds with equality, this holds with equality, this is equivalent to r hat less than r. So what's my simplified problem? I'm maximizing c equals q plus r minus k of q hat minus k of r hat, and my two constraints are that r hat has to be less than or equal to r, and this bottom one, the IC one, holds with equality.

So again, to simplify this problem, I now take this expression for what k of q hat is, and I plug this in here. So this is k of q hat equals this plus this minus this. And so when I do that, the equation becomes q plus r minus this, so minus k of q minus beta k of r, and then I get a minus 1 minus beta k of r hat. So I maximizing that subject to the constraint that r hat is less than or equal to r.

I, again, look at this expression and I recognize that the r hat only appears one way, one place, and the r hat has got a positive constant times an increasing function of r hat. So if my only constraint is that r hat equals r, I want to set r equals zero. So that's where the zero comes from. And then I'm left with a very simple problem, which is just I'm trying to maximize what's left. I'm trying to maximize q plus r minus k of q minus beta k of r.

And so then just the first order conditions say you set k prime at q equal to 1 and you set k prime at r equals to 1 over beta. And so it's the k prime at q equals 1 that makes q the efficient level. Good, so I did get that right. If the efficient level is 15,000, you ask for the efficient level at that period, but then k prime and r is 1 over beta, so you ask for discretely more money in this preference.

And so the model was that-- with the beta delta hyperbolic consumers, you didn't do any distorting. But now with beta, beta hat delta, with beta hat arbitrarily close to beta, you get this discrete jump reflecting the 1 over beta discounting versus the beta discounting. And there's just no way that you can come in as a lender and compete in this market and not exploit the consumers because this is what the consumers are seeing. The consumers are seeing this borrow 35,000, repay 25,000 contract.

They think that's what they're going to do. I can't come in and say, no, look, take a reasonable size loan, pay it out evenly, because they're going to be like, no, I would like to get the nicer car and it's a huge financial win for me to take this teaser deal. Therefore, you can't compete with that teaser deal with the regular deal, because if you're not exploiting the consumers, you can't make them that outlandish promise of you're going to get all this money and repay this little amount.

Next paper I want to do. I guess I won't get through all of these today, but. So this is a paper by Michael Grubb, "Selling to Overconfident Consumers," which has a-- it's based on an overconfidence bias. People think they know more about their future self than they do. And what he's trying to explain is one of the things I did early on in monopoly pricing is two part tariffs. Like at Disneyland, you pay $100 to get into the park and then every ride is free, or you have something where you pay some upfront cost and then you pay a low cost per unit and they extract all your surplus you're getting from price equals cost in a fixed fee up front.

What Grubb noticed is that there are actually a lot of applications where what we seem to see is not two part tariffs, but three part tariffs. So the three part tariff is this payment scheme where you have an upfront payment for the good, and then some number of units of the good sold at below cost, and then additional units beyond that sold at an above cost price, and often a very, very high price. So examples he gives. Some of them, like the credit card example, you pay some upfront fee to transfer your money to the new credit card. They'll give you six months interest free, and then after the six months are over, they're going to start charging you 19.8%.

Or you lease a car, you pay $12,000 for the three-year car lease. Car leases-- I don't know if you've done these. Car leases do have mileage limits on them. Like, because you lease the car for three years, you get 36,000 miles included free. They don't charge you for the wear and tear on the car. But then if you drive the car more than 36,000 miles, you pay an extra $0.29 a mile.

Actually, maybe for you guys, renting U-Haul trucks works this way. I don't know if you've ever rented a U-Haul truck where it's like you rent the truck and it's 99.99 for the truck, up to 150 miles of use of the truck is free. If you go over to 150 miles on the truck, it's like $0.79 a mile or something for every mile that you charge, and you run up this enormous over charge on the U-Haul trucks. Cell phones back in the day, you paid 39.99 a month. You would have two gigabytes of data free. Once you used up your two gigabytes of usage, they start charging you some extra amount of money per megabyte. That doesn't sound so big, but then it becomes an enormous charge for the overage on your data overage.

Grubb suggests that why do we see these-- we know why we would see a two part tariff, because it's efficient to price things at cost. Why would we see a three part tariff? And Grubb's argument is that we see three part tariffs because it's an exploitation of consumers who think they're better than they are at predicting their future consumption. Just like you think you're better than you are at predicting how far Buenos Aires is, you think you're better than you are at predicting how many gigabytes of data or how many miles am I going to drive the U-Haul truck or something like that.

So he gives the simple example about where the three part tariffs would come from, which is this straightforward numerical thing. So suppose firms have fixed costs of $50 of serving the consumer and then they have marginal costs of $0.05 per unit of serving the consumer. So this could be cost of providing you with a cell phone and then marginal cost per megabyte or cost of renting you the car and then cost per mile of having you drive the car.

And assume that the firm can offer a non-linear tariff t of x. Assume that consumers get benefits of $0.45 per unit for theta units, where theta is either 100, 400, or 700, and all three are equally likely. Either I get my phone and then I'm going to have some number of megabytes of data I want to use, and I don't know which of those three it's going to be, but whatever it is, I'm going to have just some inherent need for the good that I want to use those. I'm going to get $0.45 per unit using theta units, and then I'm not going to want more than theta units.

Or the U-Haul rental truck, I'm either going to want to drive it 100 miles, 400 miles, or 700 miles. I don't know which of those three it's going to be. However many I get, I get this high value for using it for those miles and then no value for using it beyond that. So anyway, all three are considered equally likely by the consumer when they purchase, and then after they purchase, they learn theta and then they choose how much to actually use the thing that they've purchased.

OK. What would happen with rational consumers? With rational consumers, we do a two part tariff. The cost of providing the units to the seller is $0.05. So you charge consumers $0.05 a unit, and then you extract all the ex-ante surplus through a fixed fee. And if consumers are going to, on average, use 400 units and they're getting $0.40 per unit of surplus, 400 times $0.40 is $160. So the rational consumer plan, you would do 160 plus $0.05. With perfectly competitive firms, you'd, again, use the $0.5 fee and you'd just be charging the 50 up front to charge cost. And this is socially efficient pricing.

What happens with the monopolist Grubb's observation on what happens with the monopolist is a monopolist in this problem would use a three part tariff. They would charge you-- instead of $160 plus $0.05 a minute, they'd charge you $180 up front. The first 400 units would be free, and then units beyond 400, they would charge you $0.45. So we just get this very high, steep three part tariff. We start high fixed fee, no cost for 400 units, and then steeply increasing price beyond 400. OK.

Why do we do this? First, why do you charge nothing for the first 400 units? Oh, sorry. I didn't say this. He's going to analyze this model where you have overconfident consumers. The starkest model of overconfidence is the consumers think they're going to need 400 units with probability one. So they don't realize that they might only need 100 or they might need 700. So they think they're going to use 400 with probability one. OK.

Why do I charge zero for the first 400 units? Why do I do 180 and then zero for the first 400 units? Why don't I do 160 up front and $0.05 for the first 400 units so you pay 180 if you use 400 but less if you don't? Well, the point is if I'm going to charge you a positive amount, like $0.05 for the first 400 units, I'm only going to charge you 160 up front instead of 180 up front. And then when you actually use 100 units, you're going to end up paying less. So you might use 100, you might use 400, you might use 700.

So if I used a plan where I charged you a per unit fee for the first 400 units, I have to give you the full discount of-- I have to give you 400 delta p back in the fixed fee. And then I give you 400 delta p back in the fixed fee, but then I only get 300 p in extra revenue because sometimes you get a better deal than you're expecting. Since you don't expect the deal, you're not going to pay me up front for that, that option of there's some chance I get to save and only pay you for 100 instead of 400.

And so because the consumers aren't willing to pay you, it's always better to charge them the full price even for the units they're not going-- this is like exploiting the consumers who think they're going to go to the gym every single day. You always want to charge them that amount up front, and then say, OK, I'm giving you this great deal where you get the 400 units for free, and then you're going to pay me in advance for all those 400 units that you're getting for free.

And then why do firms charge for $0.45 for units beyond 400? Well, the consumers are sure they're never going to use any units beyond 400. So if you offer them a discount, you can have units beyond 400 for $0.40, they're going to pay you nothing up front for that discount because they think they're not going to use the extra units.

And so if they think they're not going to use the extra units, they won't pay anything for the lower price and the extra units, so you might as well keep raising the price on the extra units until they get to $0.45. You don't go beyond 45 because then ex-post, they're not even going to buy them and they're just going to not drive the rental truck for the extra miles they need. So what we get is this three part tariff. And if you had perfectly competitive firms, perfectly competitive firms would-- again, they're forced to exploit the consumers, and so they would do something like that.

So you would give all the surplus back to the consumers with a really low fixed fee. But still, what makes it most attractive to consumers is to design the package so that they think they're going to pay the least. Which is give them the big discount up front, tell them the 400 are included, and tell them we are going to charge you these high over charges. But you think you're not going to pay them so they don't care about that. And so this is what lets the competitive sellers offer the best possible package. Yeah?

AUDIENCE: How would you tell this sort of behavior out from something like a increasing [? call ?] structure?

GLENN ELLISON: Yeah. You know, obviously, I think in some of these examples, like-- some examples that he gives are better than others on that. So for instance, the cell phone minute plan. The megabytes are system wide, not customer wide. So it can't be that you using more megabytes beyond-- there, you know that their cost per megabyte has got to be linear because it's averaged across all consumers. It's a network capacity constraint.

For ones like the car lease, you could say, is there something about the wear and tear on the car that increases as the car is driven more miles in? Is there some increasing wear and tear on the car or the rental truck because it goes more miles? Yeah, that one you would have to have something to tell apart. Again, I think the credit card one, it's going to be harder for a discontinuity in months or whatever. What Grubb does have in his paper is he has an empirical section.

We say, yeah, what did Grubb do to get this paper? Obviously the story I told you is fairly simple. How do you turn this into an *AER* paper? This is like your idea you come to me with and you explain it and it's like, OK, that's a nice, cute example. That's not a paper in the *AER.* He has a much more general model where he has-- you have a theta. The theta has some distribution.

Your consumers have some belief about theta f star. When f star differs from f, it's kind of mean preserving spread or single crossing condition. What happens? He goes out and solves the model and says, OK, this simple model I had had this three part tariff, but I do this general model and what I show is that the optimal solution is not like a three part-- not exactly a three part tariff, but it is something that looks kind of like a three part tariff. But then the second thing he does is he has this data set on cell phone users.

I don't think he says in the paper, but I think this was-- there was a time when cell phones were new, when universities provided cell phone plans to their students and you could get your cell phone not through Verizon, but through MIT. You'd get an MIT cell phone and MIT would negotiate a deal with Verizon to offer discounted cell phone service for its students. I take it that these were probably Stanford-- Doesn't say in the paper. I take it these are probably Stanford student cell phone plans. But they had four cell phone plans, one where you get this amount, just a flat fee per minute.

Small upfront charge, and then a flat fee per minute used. And then you had plans that had different numbers of included usage times like, at the time it might have been 100 minutes of usage or 400 minutes of usage or 700 minutes of usage. And so each of these plans was a three part tariff. We paid some higher upfront fee. You got the 100 minutes. And then beyond that, you started to pay these overage charges. So anyway, this unnamed university offering four plans had a structure like that. He got data on what plans people signed up for and how much usage they actually made.

So what are some of the results? So plan one customers-- let's suppose that this is plan one, one of the most popular plans. What he finds is that customers on that plan, 56% of the time they underutilized their plan and used so few minutes that they would have been better off on the per minute plan. 16% of the time they ran up overages and were on this part of the tariff, and only 28% of the time were in of sweet spot where plan one is actually better than a regular monthly plan.

And so when he estimates what the consumers would have paid had they just all bought the linear plan, he finds that, on average, the saving was 42%. Now that doesn't mean that any individual consumer made a mistake, because each individual consumer, you don't know what their belief was. But certainly if, on average, all the consumers would have saved 40% by just buying the flat rate plan, then some consumers must be irrational buying these plans and must be getting exploited by them.

And then similar calculation for the plan two customers. Again, there even more of them didn't hit this point. Again, there's a small percentage who hit the overage range, but the overage range, they pay so much more that they account for-- they're paying two and a half times as much as they would pay had they gotten the linear plan. So their overage customers are way up here. And the overages account for 23% of salary revenues.

He's got something in the paper arguing, could this be due to price discrimination rather than to exploiting consumers? And then he has some conditions on the distribution they would need for this to be price discrimination and not exploiting irrational consumers. And he argues, no, this is exploiting irrational consumers.

OK. I guess I'm out of time for today. So anyway, I have slides on here on a couple other papers. One I think is very cute is this is just another version of exploiting the hyperbolic consumers, noting that another real constraint you get with hyperbolic consumers is that you can exploit them even more if they think they're going to switch tomorrow. And so a constraint that hyperbolic consumers have is not just that I think I'm going to go to the gym more than I am, but I think I'm going to cancel.

As long as I have the option to cancel my gym membership tomorrow if I determine I don't like the gym, I can always just cancel it tomorrow so I might as well stay signed up today. And noting that this sort of friendly option of letting consumers cancel at any time can be a very effective way to get them to pay much more than they would pay if you didn't have that extra option in there. OK? And then, yeah, I think that's it. That's all I'll talk about for today.