[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So let me very briefly recap what we discussed last time. So we started off thinking about choices and risk. How do people think about economic behavior when things are uncertain, when there's risk involved? I sort of discussed and showed you sort of the main workhorse model that economists use to study risk. That's sort of the expected utility. That is a very commonly extremely useful model for many situations. So it's like very widely used.

If you ask 100 economists randomly which model explains choices [INAUDIBLE] the majority will surely tell you expected utility is what you should use. It explains a wide range of phenomena in very useful ways. You can think about lots of different things. You can, for example, think about investment behavior, finance. When you think about what should you invest in, if things become more risky, if assets are more volatile, you need to have a higher return for that, or you need to be offered a higher return to invest in those assets and so on and so forth.

There's lots of useful applications in finance using expected utility. You can think of a range of different issues. You can think about, for example, criminal behavior, about sort of the risk of getting caught. What happens when the risk of caught goes up? People engage in less crime and so on and so forth. There's lots of different behaviors that you can think about and explain using expected utility.

So I do want you to sort of take away the expected utility model is a very useful model for various applications. What we're trying to do is trying to understand are there some applications for which perhaps the expected utility model has some limitation, perhaps because of its simplicity or parsimony because there's only one parameter in there? Can we sort of alter that in some ways and try to make it more realistic in some situations?

So the key parameter of interest when you try to sort of estimate this model, trying to sort of match the data in some ways, the parameter that you'd estimate here is you need to sort of assume some functional form. This is what I did last time. I can show you this. So this here, one very commonly functional form is the CRRA utility function that's very widely used in a lot of range of settings. The feature of that is it has a constant relative risk aversion, and that has a bunch of useful properties for estimating things or making predictions.

So what you then need to do is you try to sort of estimate somebody's risk preferences. How do people behave under risk? What you need to do then sort of assume some functional form, for example, this CRRA utility function. And then the question is kind of like how do we estimate gamma the risk aversion parameter? That's the key parameter in this model.

Now, how do you do that? I showed you some different choices [INAUDIBLE]. Essentially we reveal preference. Economists believe when we reveal preference, I give you some choices that involve risk. And depending what you choose, that reveals what your gamma is. So you can do like small scale gambles, which is just like small choices between different options. Some entail more risks than others, and then you can essentially just sort of estimate using those choices, or people certainty equivalent for such choices, what's people's gamma is.

Now, what we found is that when you have small-scale gambles, people look or appear very risk-averse. They'll often decline gambles with positive expected value, which makes them appear quite risk-averse once you sort of estimate this parameter gamma. Gamma looks like gamma is above 10, above 20, above 30, really, really high.

Now, at the same time, you can look at large-scale risk. There, when you look at large-scale choices, when you sort of think about what's a reasonable gamma, people actually only appear moderately risk-averse. It doesn't look like they're particularly risk-averse. When you sort of think for those large-scale choices and we look at finance or housing or other applications that people have estimated such models, you get like gamma sort of between 0 and 2 roughly.

Now, what that then implies is if your gamma is between 0 and 2, that means essentially for small-scale gambles you should be essentially risk neutral. You should not care about really small risks that are about a dollar or two. So that sort of poses a problem because now we have sort of two contradicting answers. We have for small-scale risk, it looks like people are really risk-averse. For a large-scale gamble it looks like people are not so risk-averse. Now, we only have like one parameter, is this gamma, which is coming from the concavity of the utility function. And when we only have one parameter and sort of two contradicting pieces of evidence, we can't sort of match both, right? Because if you match one, then essentially you can match the other and vice versa.

Now, I showed you a little bit Matthew Rabin's-- what he calls the calibration theorem, which essentially is sort of calibrating, showing in a fairly compelling way that, in fact, you can formally show it's not about sort of assumptions of a specific utility function or the like. For these very sort of minimal assumptions, which is just the utility function is weakly concave, you can essentially show that declining small-scale gambles with positive expected value implies that people make absurd choices when the gambles become larger. The recitation will discuss this a bit in more detail and sort of walking you exactly, in somewhat slower speeds, through the specific example.

Now then, where we started then last time was thinking about insurance choices. This is a very nice paper by Justin Sydnor, and one very nice feature of this paper is that it involves real-world choices. So it's not like some lab experiments with some experiments. Of course, some people care a lot about undergrads. Some people might say, well, what are these undergrads choosing anyway? What does this have to do with real-world choices? I think undergrads are great.

But one might wonder, like, if you recruit people into some experiments and you see some choices, like what do these choices really reveal? Like, do we really find that these choices are predictive of real-world behaviors? So one answer to that is, well, let's find some data from the real world. Let's look at real choices that people have made in real-world settings, and this is exactly what Sydnor does.

So what he does is he has this data set of from a large home insurance provider. This is sort of 50,000 standard policies that are sort of like representative of what people choose overall. And the key outcome of interest in this study is like people's deductible choices. What's a deductible? Again, these are expenses paid out-of-pocket before the insurer pays any expenses. So you have like a deductibles of $500, you have a damage of like $200, you have to pay it all yourself. If you have a damage of $1,000, you pay the 500 and then the insurer pays the rest.

And so what he has is he has choices of a menu of four deductibles for each customer or client. So you can see both people's choice sets, and you can see people's preferred options. And that allows them to sort of say, well, if you have four options, you picked one of them. That means you preferred that one over all three others. So we can sort of essentially put some bounds on people's risk aversion.

And so we looked at this already. This is kind of what this roughly looks like. There are different deductibles, which is essentially, again, like how much how much you have to pay yourself in case of damage until the insurance payments kick in. You have the premium, which is like how much for sure you have to pay every year. And then there is the premium relative to the $1,000 policy. How much more expensive is it? That's in the third column. How much more expensive is it to choose a lower deductible relative to the $1,000 premium? And then we have people's choices which is, in this case, I guess, policyholder one shows a deductible of $250. It was a premium of $661, which is $157 more expensive than the $1,000 deductible option. OK?

And for each policyholder, the company was, in fact, sort of providing individual prices. So essentially they were looking at where do they live, what's the housing value, and so on and so forth. Sydnor knows all of that. So he knows the full set of options that people had available and their actual choices, and the options available vary by person.

How do we learn now about risk aversion? Well, so the losses to the customers are capped by the deductible, right? So any loss you have from any damage that you get, the losses are only up to the deductible. So if you have a deductible of $500, the most you can lose or have to pay in any case if any loss occurs is $500. So choosing a lower deductible then, what it does it amounts to essentially reducing that loss in case you have a damage. So if you have a deductible of $500 and decide to instead choose a deductible of $250, that means essentially in case there's a damage, in case you have to pay something, you don't have to pay 500. You have to only pay 250.

But of course, if you lower your deductible, the price of your insurance goes up, the premium goes up, and the premium you have to pay for sure. So the way you can think about this then is like, if you choose a lower deductible, for sure you have to pay more money. But in case there's some damage to you with some probability that happens, if you have some claims, you have to like pay less because your deductible is now lower. OK?

So now what info do we actually need? We need the available deductibles. Like, essentially what are the deductibles for each choice? We need the premium for each option. We need the claim probabilities and people's wealth levels because you have a utility function where there's wealth in there. I'll talk about this in one second. Any questions so far?

OK. So now one important feature in this I think was asked like last time about like, well, what about the claim rates? Well if the claim rates are really high or if people think the claim rates are really high, in some sense then, having very low deductibles makes a lot of sense because then-- and very often it happens that you have to pay. Then it makes lots of sense to have like lower deductibles.

But it turns out claims rates are actually very low. So you can see overall-- this is like the full sample, this is everybody-- people's claim rates is 4.2%. These are yearly claim rates. This is like out of 100 customers, 4.2 per year actually claim any damage. And then it varies a little bit also by choice of deductible. So there's the people who happened to choose in the end like $1,000, $500, $250, and $100. But for each of them essentially, the claim rate is below 5%. So it's very low. OK?

The second factor from this data is that reducing the deductible is very expensive. So for example, this is the full sample again. On average, purchasing the insurance where the deductible of $1,000 cost $615, we can't say very much about that choice because who knows how much the actual damages are and so on. In some sense, that's sort of irrelevant for us what that number is. What we're interested in is like what are the differences in costs of different deductibles? How much do you have to pay to lower your deductible to like $500, $250, and so on?

Now what you see here is like on average, reducing the deductible from $1,000 to $500, which is sort of what this column shows that's shown in red, costs $999.91. So if you choose then like $500, is this is a risk-averse choice or not? How do we think about that? Suppose your claim rate is like, say, 5%. Yes.

AUDIENCE: Well, I think it would be a risk-averse decision because you're paying $100 more and your deductible has gone down by $500. So claim rate for that back of the envelope calculation would need to be about 20.

PROFESSOR: Exactly. So what you're saying is you're reducing the deductible from $1,000 to $500. Now, if you think that happens with a 5% chance, on average you're going to reduce your payments by $25. So 5% times $500, which is a $25. But people are willing to pay about $100 for that. So for sure they're paying $100, and the benefit that they get is with 5% chance, at least the average customer, with 5% chance they're going to pay $500 less in case there's some damage.

That looks already pretty risk-averse. Because as Ben says, surely you're not risk-neutral, because then you would not do that. You would choose the $1,000. It looks fairly risk-averse.

Now if you go down then, if you go to like from $250 to $200, there's an additional $133.22. So that's to say reducing your deductible by another $150, from $250 to $100, makes you for sure you have to pay $133. And now if your chance is like 5% of getting like essentially a damage, that is for a 5% chance of saving $150, people are willing to pay $133 for sure. OK?

So now if you try to calibrate this, what we already know from this, the simple example is that people look extremely risk-averse. OK? So that's kind of like the exercise that Sydnor is doing, just saying, look, let's take these choices very seriously. Let's look what people have done in real-world situations. These are repeat customers, people who have done this for a long time and so on. What are people choosing? And if you sort of assume expected utility, what would people's gamma need to look like to be able to explain this data? And we can do this customer-- this is like the average rates. They can do this sort of customer by customer.

Now, what he then finds is the majority of people choose small deductibles. Lots of people choose $250, $500. Very few people choose $1,000, even among people, and this is on the x-axis, who have been at the company for 15-plus years. You would say like the first time you do this, maybe you don't understand your claim rate, you don't understand what's going on or whatever. But there's people who have been at this company for 15 years. They should kind of know at some point that claim rates are pretty low, at least on average. And so like if you have 15 years at this company, it's hard to believe that you'd still think that your claim rate is, say, about 10% or the like since it just doesn't happen very often. OK.

Now, how do we think about people choosing a deductible? Again, what you need is like the following parameter. You need to have the yearly premium. You need to have the deductible D. You're assuming no other risks to lifetime wealth, which is a bit of a distraction, but essentially you can diversify risk and so on and so forth. You also can assume that's at most one risk per year. This is, again, sort of simplification and doesn't really matter very much for the probability pi. And then for now at least, we assume accurate, subjective beliefs about the likelihood of a loss.

Now then, what is then the indirect utility function of wealth? What does the utility function look like depending on these parameters? Can somebody explain this what I'm showing here? What is this equation? Yes.

AUDIENCE: The first part says pi [INAUDIBLE] w minus P minus D. The w minus P minus D is your wealth if something happens. [INAUDIBLE] pi. And the 1 minus pi is on your utility of your wealth if nothing bad happens. [INAUDIBLE].

PROFESSOR: Exactly. So for sure, so with probability 1 minus pi, nothing happens. You have your wealth W that you had before. You have to pay the premium for sure. So in that case, you also have to pay the premium. So you're going to end up with W minus P, the premium. And then with probability pi, you also have to pay the premium, which is W minus P, but also you have to pay the deductible because some damage occurred.

All right. And then your indirect utility-- your expected utility for that year is essentially then the weighted average of these things. And pi is essentially the weight on that, which is the subjective or, in this case, assumed actual probability of a damage occurring.

Now, I sort of said the indirect utility function, utility of wealth function, what is that? What's A, an indirect utility function? And B, why is there wealth in it and not consumption? Usually we think people eat stuff and there should be like consumption here. Why do we have wealth in here? What is this? Yes.

AUDIENCE: I think they might be like assuming [INAUDIBLE]

PROFESSOR: Exactly. What is the indirect utility function?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Exactly. So usually you think what you do is if you go back to like 14.01 notes or what was done in the first I think recitation, usually what you do is you maximize consumption with several goods or one good or whatever over time, and usually there's a budget constraint and wealth is usually in your budget constraint, right? You can only consume as much as how much money you have. Could be like your income or your wealth if it's over your lifetime.

Now when you do that and maximize it, then you end up at an optimum. What you can then do is like essentially express the optimum. Assuming that you have chosen optimally, your consumption is saying you already chose whether you wanted apples or bananas or whatever. If we assume that you optimize, I can then essentially just say assuming that you're optimizing, what is your optimized utility for different levels of wealth? And usually it's a function of wealth and prices, and that's what the indirect utility function is. We can very briefly also go over that in recitation. But if you go back to your 1401 or other notes, you will see essentially that it's the outcome of a maximization problem. Usually it's like for two goods or whatever. Like, it's like income. In this case, it's wealth because it's like over-- Yeah, it's wealth, but it could be available income as well if you wanted.

Now, what the person is then going to do is like each contract gives you an expected utility in terms of how much do you expect your utility to be if you choose that specific contract. And now the maximization problem is now you choose the contract J that maximizes the expected utility or the expected indirect utility as a function of these parameters. Any questions on this?

So for each contract, we can write down what's the indirect utility function. It depends on people's wealth. So we have to make some assumption of how wealthy people are. And it depends on these other parameters. It depends on the premium, it depends on the deductible, and it depends on the subjective probability of a claim occurring in that year. We assume that there's only one claim per year. OK.

So now what we can do is then we can back out the implied risk aversion from people's choices. And in fact, what we can do is we can get upper and lower bounds on people's risk aversion from what they have chosen. Let me sort of give you an example for that.

Suppose a person chooses a $100 deductible. What does this mean? Well, this means essentially that he or she preferred the $100 deductible over all the other deductibles that were available. Right? So essentially, if you choose the 100, you get three inequalities. You get like the $100 deductible is better than the $250 deductible, the $100 is better than $500 deductible, and $100 is better than $1,000 deductibles. Now, this gives us a bound on people's risk aversion. Is it the lower or an upper bound and why? Yes.

AUDIENCE: I think it's a lower bound because $100 [INAUDIBLE] is like the lowest you can go in a lower deductible can cause more risk-aversion.

PROFESSOR: Right, exactly. Sort of like a corner solution. So like, if the person who chooses $100-- that's the person that I just showed you previously. This is the example that I showed you here. Where was it? This is a person who looks extremely risk-averse, right? This is a person essentially saying like, I'm choosing the lowest possible deductible. I'm for sure paying quite a bit of money compared to all these other options, 4% or 5% chance of not having a damage. So this person will look very risk-averse. It's the lowest possible option. So maybe the person, if there had been a $50 or zero dollars option, would have even chosen that option. We don't know because that's not available.

What you can then do is, however, you can just write down these inequalities. And so if you choose the $100 deductible compared to the $250 deductible, it will be the case that if you solve for gamma, this gives you a lower bound for or gamma. So what does that mean? We know that gamma is at least as high as the solution of this inequality will tell us, but in fact, their gamma could be even higher. We just don't know it because we don't have additional choices.

Now, if you choose the $1,000 deductible on the other hand, you'll get like an upper bound. The reasoning is exactly the same. Essentially, if you choose the $1,000 deductible, that's like the riskiest option you can choose because essentially you're choosing something you will choose not to reduce your risk in any way. So you're not willing to pay to do that. It's kind of like accepting a gamble and not choosing the safe options. But we don't know whether this person would have chosen like $1,000 on one deductible, or $2,000, or $5,000, what deductible would have been chosen because there's no other options of deductibles.

And then in between, we have essentially lower and upper bounds, because essentially we know that if you choose 500, we know that you didn't choose 1,000, and we know also that you didn't choose 250. So your gamma must be in between those options whatsoever as implied by those two options. There's a previous problem set that we posted that sort of walks you through that. We'll also go through that in recitation to go through the mechanics of that. Any questions on this? Yes.

AUDIENCE: [INAUDIBLE]

PROFESSOR: I think that's not a problem because it's just sort of flipped. So we have the gamma-- that's why you have the gamma in the denominator as well. So you get a problem if your gamma is 1, because in dividing by 1, usually people use a log utility for that. Yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yes, that's a great question. So I'll get to this in a second. So what I have done right now, and this is the typical way economists think about these things, is sort of say, let's take a model very seriously. Here's the model that I'm using to try to explain people's choices, and I'm essentially assuming everything else away. I'm essentially assuming that the person optimizes. Those no mistakes. There's no framing of facts. There's no other stuff going on, liquidity constraints and so on. And I'm taking this very seriously. I'm estimating gamma.

Now, the typical-- and this is section four of the paper, in fact-- the typical thing then that happens is other people are going to say like, well, what about if people don't understand what they're doing? What about people misperceiving the risk? What about people framing effects in terms of the way you present the choices? People like to not choose extremes, but like to choose and the middle. Can that explain the results?

There's some concerns about that. I think one thing that your explanation, for example, would be able perhaps to explain is people choosing $500 over like $1,000. It's hard for the framing effects to explain why people are choosing 250. So if you look at this figure in particular, the people who are in the company for 15 years, lots of people choose deductibles of 250.

It's a little bit harder to explain with framing effects. Why wouldn't you choose like 500 as opposed to 250? Sydnor argues that's really not what's going on. I think at the end of the day, probably it's the case that people do not want these extremes and, to some degree, I think some of what's going on in this perhaps at least contributed by some form of framing effect, maybe marketing. People who sell the insurance choices really want people-- they get paid presumably if they sell essentially these low deductibles because that's how the company makes a lot of money.

What Sydnor says there is like, well, it's actually hard to sell people on stuff that they don't like. It seems like people really seem to want these things, and maybe some of that is sort of sales pressure, but probably not everything. So I think some of what is going on is a little bit hard to rule out all of those things, but if you read the paper it's reasonable. And since it's what I'm going to show you next is like the implied gamma's like so large that even if you sort of said, OK, half of this effect is driven by other things, you would get still like really absurdly large estimates of risk aversion. But that's a great question. Yes.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Well, to some degree, in some sense, I think the way they sort of-- The question was like, is the company deliberately giving people choices that leads them to choose low deductibles? And therefore, do we sort of overestimate people's gamma? To some degree, yes, but it's not like just we have low deductibles available. There's a $1,000 option available. I think maybe what you're alluding to is like there's some sales pressure and so on going on. That may well be true, and people sort of might emphasize risk and make it particularly salient and make customers nervous and say like, look, these floods and so on are going on, and really, low deductibles are good. I think to some degree, that's true, but you have to be pretty compelling in your reasoning.

One other comment is like, there's lots of other examples of people choosing low deductibles and sort of extended warranties and so on. And in lots of cases, for example, if you look at like iPhones or iPads and so on and so forth, laptops, et cetera, Apple in particular, but other companies try to sell extended warranties that, if you actually did this exact same calculation, are not worth engaging in or that reveal essentially extreme risk aversion among customers.

I myself am looking at this kind of research. There's lots of research that argues that people shouldn't choose extended warranties. Of course, then when I bought a laptop the last time, I was like, of course I don't need extended warranties. And then, of course, soon after, the laptop broke and so on, and I didn't have a warranty and so on. So of course, in specific examples that may happen, but on average it's not a good idea to do. OK.

So now what does Sydnor find? So now here in this table, you can see the implied estimates of gamma. Remember, what we said is what we think is a reasonable gamma is somewhere between 0 and 2 for large-scale choices. He has different types of assumptions. He has like lower bounds and upper bounds of gammas. And what essentially you see is the gammas are like, depending what the assumptions are, in the hundreds or in the thousands. It depends a little bit what people's wealth is. So essentially depending on how much you assume people-- So the data that he does not have is what is people's wealth actually like. How much money do people actually have?

So he makes some reasonable assumptions by saying, look, these are people whose houses are like $200,000, $300,000, $400,000 worth. So presumably, and we know from other data sets roughly how much money people have available. So then depending on essentially what utility function you assume, he is mostly using like CRRA utility, like then is wealth of like $1,000,000. You can look at like $100,000, $50,000, $5,000, and so on. You can also use CARA utility, which is constant absolute risk aversion, and so on.

And essentially, what you have to assume is like extremely low levels of wealth or something like $5,000 to get into like sort of single digits or double digits of gamma. It's extremely hard to sort of get estimates that are reasonable in the sense that we think are actually reasonable parameters of gamma. Any questions on this?

OK. So now, why do people choose those small deductibles? And this is what you were saying before. So one classical explanation would be, well, they must be really, really risk-averse. Their gamma must be really high. There you might sort of say, well, we know already that gamma shouldn't be that high from some other choices. So that's sort of hard to reconcile. Was there a question? No.

Second, you could say, well, there's really high objective probability of claims. We know that the objectives for probabilities are only 4% or 5%. You'd have to sort of have probabilities of claims that are like 20%, 30%, 40% to be able to match these data. Now, it be that people have risk misperception. It could be that people really think the probability is actually 20% when it's, at the end, only 5 or 4.

Now, what's sort of inconsistent with that is that like repeat customers, people that have been at the company for 10, 15, 20 years, are making very similar choices. In some sense, it's hard to reconcile that everybody is misperceiving this risk year after year after year and spending lots of money. There's some questions on is it borrowing constraint? Is it like people just don't have enough money? Are they worried about having to pay these deductibles? That also seems quite unlikely because in fact, the deductibles are particularly-- we're not talking about like $5,000. People could sort of, if they really faced these borrowing constraints, they could save and so on. So Sydnor also argues that that's not going on.

There's some questions about marketing social pressure and so on, which is-- of course, the company has very much like incentives to sell people these kinds of deductibles. I think some of that is probably going on, and it's hard to rule out entirely. Again, like, it's hard to actually sell people stuff they don't really want to. So there must be lots of social pressure to, in fact, do that.

Menu effects. Already talked about it a little bit. There we think that maybe menu effects can explain why people choose to 500 or, like the interior choices, why they don't choose 1,000. It's hard to explain with menu effects why people choose 50 over 500.

So then Sydnor's preferred explanation is then reference-dependent preferences and loss aversion, which we're going to talk about next. Yes.

AUDIENCE: Is there any data on whether the probability of the claim is correlated with which menu option the people chose?

PROFESSOR: Yes, there is. So you have it here. So it's sort of weakly correlated in a sense. Like, if you choose-- or weekly negatively correlated. So people who choose $1,000 have 2.5% claim rates, and the $100 are 4.7. So that's kind of not quite explaining things either. Yeah.

AUDIENCE: I think the use of [INAUDIBLE] regression to control for fact that those with lower deductibles may claim multiple times.

PROFESSOR: I see. Yes. Great, yes. OK. So what do we learn from this? I think this is very much sort of confirming some of the lab evidence data on relatively small-scale gambles. These are not like small-scale gambles of like a dollar or 2 or 5 or 10. These are about several hundreds of dollars, but they're not about hundreds of thousands of dollars. So these are relatively small relative to like people's lifetime wealth.

And what we see for those kinds of choices is it really looks like people appear to be very, very risk-averse. This is what I was saying already previously. When looking at sort of reasonably small-scale choices-- and I count the Sydnor evidence as reasonably small-scale choices because it's not about like hundreds of thousands of dollars-- you essentially see that such choices imply people seem very risk-averse. These choices imply enormous risk aversion for large-scale risks. But people are not avoiding all sorts of large-scale risks. People take on lots of large-scale risk. So in some sense, that can't be true in some ways.

We also find that individuals are moderately risk-averse a large-scale risk. People are taking on some risk, as I said. Now, if you sort of take that seriously and say, well, people might not be risk-averse, that in turn implies that people are nearly risk-neutral. They should be nearly risk-neutral for small-scale risk. So it can't be that both of these things are true at the same time.

Now, in fact, there's a much older paper that's a very famous and seminal paper by Kahneman and Tversky. Kahneman got the Nobel Prize, in fact, for this work and similar work. Kahneman is a psychologist and they were doing psychology experiments, and it just turned out that these psychology experiments they're extremely influential in affecting how economists think about risk and risk preferences, and in particular, sort of preference-dependent preferences.

And so what Kahneman and Tversky were doing at the time way before a lot of this other literature that I just showed you-- A, they showed even more sort of evidence against the expected utility model. But perhaps more importantly, they also proposed an alternative model in saying, look, here's some choices that we think are hard to explain and hard to rationalize using expected utility. Now here's a different model that can explain things perhaps better in some situations.

So what did Kahneman and Tversky actually do? The experiments are actually extremely simple in various ways. They're very clever and very clean, but actually very simple. And so what do these experiments look like? These are survey responses. Like, essentially they just asked people about what would you do in different situations. This is not what economists were doing at the time. Economists at the time were saying like, revealed preferences are important. I need to make you do actual choices. Whatever you say in surveys doesn't matter because who knows whether you actually mean what you say.

It turns out that the survey responses are actually-- these are hypothetical stakes, but it turns out that if you do this with actual stakes, you find very similar results. And the experiments were as follows. There were things like questions like, which of the following would you prefer, kind of like I showed you in the first class at the end of the survey that you did. Would you prefer option A, which is a 50% of winning $1,000 or 50% chance of winning nothing, versus option B is like $450 for sure? And so they did a series of these types of questions.

Now, one of the things that then they showed is one key prediction of expected utility is that, as I said before, people only care about final outcomes and their associated probabilities. Kahneman and Tversky show a bunch of different striking contradictions of that. I want you to focus on the first row-- problem three and problem three prime. And look at that and tell me what of that example is contradicting expected utility.

I don't know if you can see this. So let me just read this for you if it's hard to see, but think of them for a second. Problem three says, "Would you prefer an 80% chance of $4,000 over $3,000 for sure?" What do you see below then? These are always 100 people. You see the number of people who preferred one option over the other. So 80 people preferred the 3,000 option. 20 people said I'd rather have the 80% chance of $4,000. That's problem three.

Problem three prime is about what they call negative prospects. Would you prefer an 80% chance of losing $4,000, or for sure losing $3,000? And there you see 92% choose the first option and 80% choose the $3,000 loss for sure. Let's start with the left example. What did we learn from the left example? Are people risk-averse, risk-neutral, risk-loving, or what have we learned from that?

AUDIENCE: Risk-averse.

PROFESSOR: Risk-averse, and why is that?

AUDIENCE: If you calculate the expected monetary value, I guess you would expect to [INAUDIBLE].

PROFESSOR: Right. So we said before somebody is risk neutral if the person is indifferent between two options when they have the same expected monetary value. And so in this case, an 80% chance of $4,000 is $3,200 in expectation, right? But there's a bunch of people who say, I'd rather get the $3,000, which is less than 3,200, for sure. Which means essentially they take a surer-- like an option that's for sure they get. They prefer that over the uncertainty of getting 0 versus 4,000, which on average gets them more, 3,200.

Expected utility would say if you choose that option, it must be that you are risk-averse. Now, not everybody seems risk-averse. There's 80 people out of 100 choose that, or I think it's 80% choose that and the remaining ones choose the other option. These other people we don't know much about. They're not very risk-averse, but they could be also risk-averse just like it was revealed in their choice there. OK. So from the left side, we say we know at least 80 people, or 80% of the sample, is risk-averse. Now let's look at the right side. What do you see on the right side? Yes.

AUDIENCE: Well, they seem to be risk [INAUDIBLE], because unexpected risk minus 3,200 versus minus 3,000, but they prefer the minus 3,200 on expectation.

PROFESSOR: Right. So if you look at the left option, it's minus 3,200. So the left option has more risk, and it also has a lower expected monetary value. As you said in the expected monetary value, it's just a flip of the problem three. It's minus 3,200, and the other one is minus 3,000. So if you're in a risk-neutral, surely you would choose the minus 3,000. If instead you choose the minus 3,200, well, it must be that you really appreciate that there is additional risk there. So it looks like you're risk-loving. OK.

And now we find here that 92% of the sample choose the first option, but at the same time, with the same people, when they're offered or given the choice between problem three, 80% of people choose the other option. So there are a bunch of people essentially that choose the $3,000 for sure in problem three. At the same time, they choose the minus 4,000 with an 80% chance when given problem three prime.

So that means essentially we have two choices here. One choice says people are risk-averse. The other choice says people are risk-loving. In expected utility terms in our world, this cannot happen. The reason being that we have one parameter, which is gamma. Gamma tells us how risk-averse or how risk-loving or the risk-neutral alike. That parameter tells us everything about your risk preferences for all choices that I'm giving you. It cannot be that you're simultaneously risk-loving and risk-averse. So this evidence essentially sort of rejecting the expected utility model or cannot explain this behavior. Does this make sense? I sort of wrote this down here, but I think I said everything that's to be said. Any questions on this?

So where we're going to go is essentially A, these people seem to be behaving differently for gains versus losses and, in addition to that, so not only do people seem to dislike losses-- I'll show you some evidence of that-- but in addition, people seem to be risk-loving for losses and risk-averse for gains. And that's kind of what Kahneman and Tversky are claiming. And once you make that claim, you'll be able to explain these patterns in the data.

Now, a second thing that they show is the following. They show in problem 11 and problem 12. I'll let you read it for yourself. And essentially, you see that 84%-- so the choice here is in addition to whatever you have, for sure you're going to be given $1,000, or shekels I guess. You are now asked to choose between 1,000 with a 50% chance and 500 for sure. OK? 84% say option B.

Problem two is, in addition to whatever your own, you have been given 2,000. You're now asked to choose between option C, which is minus 1,000 with 50% chance, and option D, which is minus 500. 69% of people here say they choose option C. OK? So what's the problem with this? Yes.

AUDIENCE: I mean, at the end of it, you would say that problem 11 option B, you end up with 1,500. Whereas problem 12 option B, you still end up with 1,500. People were inconsistent with these decisions based on how the question's framed.

PROFESSOR: Exactly. So framing matters, or reference to points matter. And so what assumption of the expected utility model is that rejecting?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Exactly. So we postulated that only final outcomes matter. Here, if you write this down, it turns out option A and option C are exactly the same in terms of final outcomes. Turns out option B and D are also exactly the same. I'll let you look at this for a second. So A and C are the same in terms of the final outcomes. B and D are also exactly the same in terms of final outcomes. So that means essentially when you compare A versus B and C versus D, If you only cared about final outcomes, you cannot choose different things here in this lottery. It cannot be that your utility is defined over just final outcomes, because you just told me you like two different things for the exact same thing. And there's 84% who choose option B while 69% choose option C. So there's a bunch of people who essentially switch. Is that clear? OK.

So now, what is that? What's going on here? Well, where we're going to go is like well, it seems like, exactly as you say, framing matters. How you frame the problem matters, but why does it matter? Well, it's because we setting like a reference point. In the first place, in the first example, the reference point is like 1,000 and there is like a gain relative to 1,000. We look at like people evaluate this gamble as like gains.

In the second example, the reference point is set to like 2,000, and now people essentially evaluate this gamble as losses relative to the 2,000. And of depending on how people think about gains versus losses, it turns out that people make different choices. I'll be more formal about this, but that's essentially the idea of, like in Kahneman and Tversky, their prospect theory they proposed, why can explain these behaviors as opposed to expected utility? OK.

So now what are sort of the most important points? There's three of them. I'm going to show you two and then get to the third one at the end. So one is like what matters a lot for people's behaviors is changes rather than levels. So what they argue is utility is not defined by people's final status-- how much they end up with-- but as opposed to like how much changes relative to some reference point. That could be changes relative to like a status quo. How much do I have right now and how much does it change positively or negatively? Or, it could be relative to some expectation or some other reference point.

What I showed you here in this example is like this is not the status quo that people evaluate their utility against because the outcome is always the same. The reference point here is kind of like the expectation I sent you and say, OK, here's 1,000. And now what seems to be happening is that people's expectations, in fact, become 1,000. And then relative to that expectation, people are going to evaluate gains and losses. And similarly, if I give you-- like, say you get 2,000, again, people will evaluate the utility or the outcomes as gains and losses relative to that.

Second, there seems to be loss aversion. Losses loom larger than gains. That is to say people dislike a loss that's as large as a gain by a lot more than they like the gain. OK? So it's like if you say you lose 100 versus you gain 100, people really dislike losing 100 a lot more than they like gaining 100. And once you sort of postulate that, well then, that can explain why you reject a bunch of gambles. Because if say if the gamble is like minus 10 with 50% chance and plus 11 with 50% chance, well, if you dislike the losses a lot more, if you dislike losing $10 a lot more than gaining $10 or $11, well then, you're going to reject this gamble even if the expected value is positive. OK? We'll write down sort of like a utility function next time. That's sort of does this more formally, but essentially those are the two key areas. There's a third one, which I'm going to get to in a second.

Now, the key part here is like there is essentially reference dependence. People evaluate their outcomes relative to some reference point. What kinds of examples do we actually have a reference point? When you look in the world-- and again, that's kind of very much what I'd like you to do-- when you think about things that you see in the world, what kinds of examples do we have that people care about reference points as opposed to final outcomes? Yeah.

AUDIENCE: My first thought was, you know how we have to pay $0.10 for shopping bags? Where I was from a couple years ago, we did it so that if you used a reusable bag, you got $0.05 back. That was like not effective. [INAUDIBLE]

PROFESSOR: Right. Yeah. So exactly. When you think about pricing of different options, it matters a lot to people. It seems that when you add $0.05, were to subtract $0.05 or $0.10 or whatever, the same changes about like it's just $0.05 at the end of the day that you have to pay more or less, depending on whether you use a bag, your choice of using a bag versus not should not be affected by whether it's framed or put as like a loss versus a gain or like a whatever rebate or whatever that you might get. And I think that's true for many different-- that's true for shopping plastic bags or any other shopping bags, but it's also true for many other options. It depends a lot whether you sort of add on that option or whether you get sort of essentially subtracted the option as some discount. Yes.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Right. I think some of it is not exactly-- I think there's some legal constraints to that specific behavior because you're not supposed to sort of trick people, but surely some of that is going on. People love discounts. Like, they love making good deals. They love sort of getting things cheaper in various ways compared to some reference point. And the reference point, since it's hard to tell what actually should the price be, is often like the previous price. It's something that seems really expensive and now they're getting it for less money. And so if you get it for half the price, even if it's still quite expensive, you think like you saved half of the price somehow.

AUDIENCE: So how you feel about what you have may depend on what you see your neighbor having.

PROFESSOR: Right. So reference points could not just be prices or sort of things that affect yourself in certain ways, and price is the key example overall. The reference point could just be your environment, your social environment. We kind of talk about this to some degree also when we talk about social preferences. People care a lot about what others do, and their reference point might very much be formed by what their neighbors do. Like, how big is their house? How big is their car? Do they have a swimming pool or not? It could be also like your neighbor's-- like, how good are they at school? Are they smart, how good looking? And so on and so forth. So when people evaluate certain outcomes, they often don't evaluate the levels, but rather kind of how much another person makes or whatever other people's outcomes are.

And part of the reason might be that like evaluating absolute outcomes is really hard. It's very hard to actually understand. Is $10,000, $50,000, $100,000, a million-- how much money do you actually need to be happy or that you like, or what kinds of outcomes are you excited about and so on? But it's much easier to say you have something and another person does not or the other way around. It's much easier to compare with others. It's very hard to actually evaluate absolute outcomes, because who knows how you should feel about that. Much easier to say I have this, they don't, or the other way around. Yeah.

AUDIENCE: People sometimes don't want to let go of fallen stock they own because they don't want to see that they sold at a good price.

PROFESSOR: Right. Exactly. That's called what people call the disposition effect. There's actually pretty large literature studying this and lots of debates on whether it's really going on or not, how important it is, and so on and so forth. But sort of one very basic stylized fact is that when people are looking at stocks that they might want to sell, they are much more likely to sell winners compared to losers. Now, why is that bad, or should you not do that?

AUDIENCE: Is that because the winners [INAUDIBLE]

PROFESSOR: There's a bit of a question of is there like momentum or reversal or the like. But if you believe in efficient markets, which many economists do, or at least some of them who are in Chicago-- if you believe in efficient markets, then the current stock price should essentially not be-- or like the previous losses and so on should not be informative of what's going on in the future. So in expectations, your losers and your winners, if there's any information about the future, about the future valuation, that should already be incorporated in the price. So if you look at two stocks, one lost some money and one gained some money, they're equally likely to make you money. And so if anything, what they show in these papers is like there's momentum. This is what you're saying is that the winners actually are, in fact, going to be more likely to try to increase value compared to the losers. But people tend to essentially want to sort of realize gains. People seem to be happy about making gains. They seem to be very reluctant to sell the losers, and there's some questions on how costly is that actually, but that's a very robust pattern in the data.

The same is also true-- I'm going to show you on Monday, the same is also true for houses. People are much less likely to sell their house when it has lost value compared to when it gained value. I'm controlling for a bunch of things. Yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah, so that's a little tricky. There might be other things going on, but exactly. It could just be that like if you don't go to a movie when you have bought a ticket yourself, it's sort of perceived as a loss and you really don't like that. It's a little complicated to think about this in simple terms, because in some sense, the loss is still there. But anyway, yeah, exactly. It's like you lost money. It's almost like as if you lost the movie ticket and didn't get some value from it. There's a bit of a question whether you sort of integrate these two things or not because people think about monetary terms versus other things often in separation, but I think some of that is exactly. People feel like they have a loss if they don't take advantage of the movie tickets.

So I'm going to show you some examples that are actually much more basic in some ways. When you think about like visual illusions, this is what's called like the size contrast illusion. So one of those things is like when you look at circles or things that's supposed to look the same, and in fact are the same size, they look quite different. Depending on what you contrast things with, things look quite different. For example, if you look at these circles in the middle, they're in fact exactly the same size, but in fact, they don't really look like that. Similarly, if you look at these two circles, it's perhaps more stark. These circles are, in fact, exactly the same size that are in the middle.

Every time I'm teaching this class, I sort of have to convince myself that they are actually the same size. So I print it out and sort of measure to make sure that it's actually the same size, and they are. I checked. But when you see this, even if you know it's an illusion, it's very hard to convince yourself that it's not. But they are the same size.

When you look at these bars, when you look at the upper black bar and the lower black bar, again, it seems like somehow the upper one is like wider, but in fact, it's not, and again, it's all about the contrast to the other side. There's this one which throws me off the most. When you look at fields A and B, they are the same color. That's hard to believe once you see it. They are actually the same color, and again, I would print it out and actually put them next to each other. They are the exact same color. I checked.

There's also a video that you can watch that sort of shows this. And exactly what's happening is the way we perceive colors is coming from contrast, right? So like black does not or gray doesn't look as gray when white or darker gray is next to it. Similarly here, if you look at the gray of this bar, this is the exact same gray. Again, you can print it out and sort of look at it. It's the exact same. When you look at it, it just does look like on the right side it's darker than on the left.

Now, there's plenty of examples of reference dependence for vision. There's tons of them. They're kind of quite interesting. At some level, it tells us something about the brain. Like, in some ways, we can learn about essentially the way we evaluate certain outcomes overall is essentially-- it's much easier for us, or we think about in contrast, it's very hard for us often to think about levels. But of course, that's vision. So in some sense, what did we learn really about utility?

One thing you can look at is bronze and silver medal winners at the Olympics. Presumably or arguably, winning a silver medal is better than a bronze medal winner. If you look at these two women, one of them won the silver medal, one of them won bronze. So what people have done and psychologists have done is like took actually a bunch of pictures from like ceremonies of bronze and silver winners and just looked at who looks happier. And when you do that, essentially the bronze medalists look on average happier than the silver medalists. Presumably it's because they just missed gold while the other person sort of won bronze because it's great to be third as opposed to fourth.

And there's more. You can read more about this, but that's sort of like one finding. Now, it's true for lots of different feelings, perceptions, judgment, and so on. People compare stimuli in various ways when it comes to temperature, when it comes to all sorts of things relative to reference levels. It's very hard for them to evaluate it in absolute terms. So it's easy to say-- when you look at water, it's very hard to say what's the temperature even if you sort of practice it a lot. It's very easy to say one bucket of water is warmer than another. It's very hard to sort of say it's like 70 degrees or 60 or whatever. It's very hard to sort of understand what absolute temperatures are, and I think that's the same in some ways for a lot of consumption or other decisions. It's very hard to say how much would you pay for this or how happy should I be with certain outcomes, because people need some reference, and often then the reference are either some expectations, or their neighbors, or just comparing between different options.

So that's the example I mentioned previously, which is it's much easier to compare your income or any outcomes-- your grades, et cetera-- compared to what your friend has. It's much harder to say how much an extra $1,000 or having $50,000 per year, is that good or bad? It depends a lot on what you compare it against.

So what people tend to do, and this is exactly what Kahneman and Tversky were postulating, is that people compare the outcomes relative to reference points. And again, we're going to write this down in more detail at the beginning of next class, but what Kahneman and Tversky were postulating were essentially two things. They're postulating A, there's a reference level of consumption or any outcomes. We can talk a little bit about what actually is this reference level, where it's coming from. For now, we just assume there's a reference level that people compare the outcomes against.

And then against that, outcomes are compared. And in particular, the function is like steeper on the left compared to the right, right? Essentially losses loom larger than gains. So like going down by 10 units on the left or 1 unit on the left is more painful than going up on the right relative to those reference points. OK? And so that's all said.

So then what experimental evidence do we in fact have for that? So one example-- so we have some experimental evidence of this. And I'm going to show you next time a bunch of different examples of real choices, starting from golfing, to selling houses, to lots of other outcomes. But sort of the experimental evidence, the earlier experimental evidence that people showed were preferences over risky gambles. I showed you some of those already, and then in particular unwillingness to trade different options compared to like an alternative option that I showed. That's what people refer to as the endowment effect.

So let me show you first the gambles. We already had that in some sense. These are like essentially gambles. People seem really risk averse when they are offered these gambles. Kahneman and Tversky would say people are essentially loss averse. People really dislike the loss of $10 relative to the gain of $11. We can explain that behavior. And that's a very robust finding that people decline these kinds of types of gambles. Kahneman and Tversky would say that's evidence of loss aversion.

Now, you might say these are really small gambles. Do we really care about them? Well, once you do this with $500, $550 or whatever, large amounts, people also do the same. It's hard to do this with like real-world money because it's quite a bit of money. It turns out there's actually some studies who do that as well. They did this for real with MBA students, financial analysts, and rich investors. And even those people like tend to then turn down those kinds of gambles. So there seems to be quite evidence of lots of loss aversion. Given the choices between gains and losses, people seem to really dislike losses.

Now, what's the endowment effect? That's sort of the perhaps most famous evidence in this domain. It essentially is people, when endowed with a certain item, they really do not like to trade. And so the way you would do an experiment-- you would give people an item, and then you ask them essentially-- and a randomly selected fraction of people are given an item. And then either people are offered the choice-- you would you either keep your item, or would you get another item that's sort of the same value overall? Or, what people do is like they have some experiments-- there's many of these experiments-- you give some people one item and some other people another item, and then see would you like to trade with each other. And what you see is essentially people are extremely reluctant to trade. There's many of these kinds of examples.

Like, for example, so one example-- sorry, I skipped that. One example is like if you just give people an item-- usually it's like a mug-- if you give people a mug and asked like, how much do I have to pay you to give me to sell me this mug? People say large amounts. They would say like $5. If, instead, you asked them, here's a mug. Would you like to purchase this mug? And sort of controlling for how much money they have and so on. If you ask them like, would you like to buy this mug, people would say like $2 or something. It's the same mug, and essentially their choice depends on whether you give them the mug versus whether you endow them with the mug, which is where the endowment effect name comes from, or whether you just sort of like ask their willingness to pay when they don't have it. There's tons of different experiments that do exactly that.

Now, a different version of that is like so this one is about like buying a mug, selling a mug, buying and selling prices. Different versions of that is to say you have a population of students or different people, and what you do is essentially you find two items that on average have the same valuation. Knetsch was doing that. And so he had like mugs and pens, and he sort of calibrated such that on average, when you just asked, people's willingness to pay for these mugs and pens are like on average roughly the same. Of course, there's variation, but on average it's the same.

And then you offered half the students mugs and half the students pens and then offered in exchange. He also offered in exchange for an addition of $0.05 and saying like, OK, maybe you're just exactly indifferent. So I'm giving you $0.05 in case you exchange. And it turns out that the mug people like to keep their mugs, and the pen people like to keep the pens, and 90% of people do that. That's one of the most robust findings in experimental economics. There are some questions on expectations. Do people expect to keep the mugs and so on and so forth? There's some complications for that, but the basic result is very robust and has been shown in many different settings. Any questions on this?

OK. So what's going on here? Well, essentially, the reference point seems to be essentially affected by ownership. So if you own a mug, your reference point is owning a mug. So now if I asked you, would you like to sell me this mug? Well, now you have a loss of a mug. So essentially, you're in the lost domain of mugs. So I have to pay you more money to compensate you for that.

If, in contrast, you do not own the mug, so like you have zero mugs, your reference point is zero mugs, and I'm asking you, would you like to receive a mug or like buy a mug from me, then you're essentially on the gain domain. And sort of what I showed you here previously-- one second. So when I'm asking you to purchase your mug, essentially you're on the left side of this figure. You're in the loss domain. Your marginal utility of mugs is very high. I have to pay you a lot of money to receive the mug from you.

In contrast, if you have zero mugs, I'm asking you, would you like to purchase a mug, you're on the right side of the suffering point. So essentially now you're in the gain domain. Now you're not willing to pay a lot of money because you're gaining a mug. On top of that, you also on the gain loss domain of money, but like I'm setting that sort of aside. Does that make any sense? OK. And so people hate losses more than they like the gain. So they stick with the mug. And similarly, that's the same thing for the pen owners.

There's lots of different examples of those kinds of behaviors. For example, law school students were asked to assess compensation for pain and suffering damages in one study. This is expected to last-- or in this example-- is expected to last three years and be quite unpleasant. There's no impact on earnings capacity. For example, it would be extreme stiffness in the upper back and neck. I think that would probably affect your earning capacity, but anyway let's assume that's not the case.

So then some students are led to imagine they were being injured. How much would you be willing to pay to get better? OK. So your reference point is like you're injured. How much are you willing to pay to get better? And people said 151,448 on average. Now, another group of students was-- this is randomized-- was led to imagine being uninjured. Like, how much would I need to pay you to accept the injury? You're not injured. Now I'm asking you, how much are you willing to pay, or how much do I have to pay you, how much do I have to compensate you to accept this injury?

Now, you would think this is like the same thing, but what's your price of not being injured? The price of health should be independent on which way I'm asking you. Turns out, when people have their good health, they're willing to pay a lot more to keep it or to not lose it compared to like when they have lost it and they want to regain it.

There's another quite nice example. I'm showing you here. One second. So in case you didn't see, that was Dan Ariely, who has several nice books. One of them is predictably irrational, which is a quite nice read. So there's lots of these kinds of examples of loss aversion or what's called the endowment effect. When people have stuff essentially, they're not willing to part it. When they don't have it, they're willing to pay less.

One nice thing that Ariely was also talking about-- a little bit about how people then explain these kinds of behaviors. We're going to get back to that in I think lecture 20 or something, which is about kind of like I think what's going on here is very obvious. Potentially, people are loss-averse and they're randomized into like gains or losses, and that's kind of what's explaining their behavior. But people don't necessarily understand that they're being randomized in one of the other conditions, and they then sort of try to explain their behavior in various ways and saying, I want to tell my grandchildren about this. But in some sense, they're sort of like rationalizing their preferences and they don't necessarily sort of understand these stories, and they don't necessarily understand where these preferences are coming from even though we kind of know, because the person has been manipulated into those kinds of choices. We'll get back to that.

Now, the third part of Kahneman and Tversky's prospect theory of their paper is essentially what's called diminishing sensitivity. I'm going to get back to this and sort of write this down in more detail on Monday, but essentially it's like people are risk-averse in the gain domain, but risk-loving and the loss region. What does it look like? Essentially the utility function is a little bit different from what I showed you before. It's not only like more steep on the left relative to the reference point compared to the right, but it's also concave on the right and convex on the left. And what that gets you and what that buys you essentially is that people are risk averse in the gain domain, but people are risk-loving in the loss domain, and that sort of is needed to be able to explain some of the behaviors that I showed you early on in the Kahneman and Tversky evidence.

So again, I'll tell you about this in more detail and sort of write this down more precisely. And then we're going to talk about like many different applications. And particularly, there's the endowment effect which is kind of what I talked to you already about before. There's labor supply, employment, and effort depending on what people expect, how much they earn. They make different choices about how many hours they actually work depending on the expectation, whether it's above or below a reference point. People are reluctant to sell their houses if they're at a loss compared to again even for very similar houses. And marathon-running people are trying to reach essentially certain targets. People like to be below four hours or three hours and the like.

There's the disposition effect that we mentioned, which essentially is people like to sell winners compared to losers. And the insurance choice that I already showed you before-- there's some evidence on violence in the household. The particular example is about football or games essentially-- I think football games where essentially depending on whether people expected a loss or win of their team and when then a loss or a win of the team actually happens, there's more violence when there are unexpected losses compared to expected losses which is, again, consistent with loss aversion.

What's next? So next we're going to talk about many applications of reference dependence on Monday. I'd like you to sort of read the Kahneman and Tversky paper, at least the first pages to get some sense of their work. And then on Wednesday, we're going to start talking about social preferences. In particular, we can have some experiments in class. There's going to be no readings for that. There's the opportunity to make money in class, at least some of it.