

[SQUEAKING]

[RUSTLING]

[CLICKING]

**FRANK SCHILBACH:** Welcome to Lecture 17. This lecture talks about state-dependent preferences, projection and attribution bias. The plan for today is to talk about preference changes. So we're going to look at different circumstances in which people's preferences, predictably and sometimes unpredictably, change over time. That is to say, sometimes people are hungry, tired or in pain, and their preferences for certain outcomes may be determined by that underlying state of like hunger, for example.

Now, it's pretty clear that people's preferences change according to those states, as in like when somebody is hungry, they want different things. They want to eat different kinds of things than when they're not hungry and so on. But people nonetheless have trouble predicting their preference changes. That is to say, if somebody is hungry, they may have a hard time predicting how it feels when they're not hungry and vice versa.

So notice that this lecture is sort of a continuation or natural continuation of our lectures on beliefs from before. So before, we looked at people trying to learn about unknown information around them, and we looked at various deviations why people are not able to learn. Now we're going to look at people trying to learn or failing to learn at times about their own preferences.

And one particular case here is what's called production bias, which is people's lack of or inadequate ability to predict their own preferences, in particular for different states of the world in the future. I'm going to be more specific about that, but mostly, we're going to talk about projection bias and the paper by Loewenstein et al., which is in the readings. We're going to talk very briefly about attribution bias. That's going to be covered in recitation.

OK, so although not typically emphasized in economics, the simple and obvious fact is that our preferences change over time. There are ways in which our preferences change. One is short-term temporary fluctuations, and you can call them sort of state-dependent preferences. You can call all of what we're going to talk about state-dependent preferences. Here, the state is a short-run physiological or psychological state, for example, hunger, pain or the like, or a psychological state such as mood. We're going to talk about all of these in more detail.

Second, there are long-run systematic changes. These could be due to own choices such as addiction. So like if somebody has been drinking alcohol for the last 10 years, their utility from drinking alcohol or not drinking alcohol is quite different from somebody who has not been drinking ever. The same is true for smoking and the like.

Or it could be independent of one's choices, for example, aging. So predictably, people's preferences change over time when they're 20 years old versus when they're eight years old. Notice that for addiction, there's also some short-run temporary fluctuation such as when people have just smoked a cigarette versus not. Their preferences for an additional cigarette might be quite different.

But then there's also long-run systematic changes, which is like if somebody has been smoking a lot during the last 10 years versus not. So smoking or addiction in particular has both short-term and long-term temporary and systematic changes in people's preferences. And then there's adaptation to changes, which happens often for big changes such as standard of living or small changes such as mug ownership.

What's an example here? For example, if people win the lottery, they become happier. They become actually a lot happier very quickly, but then that sort of reverts back to like their previous standard of time. People, a lottery winner seen like a few days later, clearly are happier than people who have not won the lottery, but their increase in happiness decreases over time.

So these are three different types of sets of preference changes. You're going to mostly focus in the lecture today on, number one, short-term, temporary fluctuations. Now, what am I talking about here? So one example here is hunger. Now, studying hunger is difficult in experiments because it's unethical to starve people. And you could, of course, find hungry people and give them food.

That's in some sense what was done once in what's called the Minnesota starvation experiment. This is an experiment that was motivated by the fact that there was lots of soldiers in the Second World War from the US and Europe and in many other places, and these soldiers had been starved from the war for a long time.

And they came back from the war to come back to the US, and one simple question was, well, when you have people who have starved, have not eaten for a long time or very little for a long time, how do you best integrate them into society, or how do you best feed them. Do you feed them very quickly, like a lot, or do you sort of gradually increase their food take over time?

So then the Army did some experiments with healthy volunteers such as this guy here on the left. And people were starved, literally starved by getting very little food for a while, such that, I think after a few weeks, the guy would look like this. And then the Army would look at people's behavior over time. They were mostly interested in like once you have somebody who has not eaten for quite a while, once you give them more food, what's happening to their behavior and what's the best way of doing that.

Coincidentally, they also reported what was happening in the first part of the experiment when people were starved to start with. And so as expected, people get extremely focused on food when they don't eat for a while. They get really, really interested in learning about what's going on with food and when they get next food and so on. But as it happens, it seems to be that their preferences also change in various other ways.

In particular, people seem to lose interest in lots of other activities. It is reported decreased alertness, lack of self-control, general apathy. People are just sort of like almost like a different person when they are almost starved compared to when they have eaten a lot. Here's some testimonies from these soldiers. One of them is, the acquisition of food-related items was a reasonable extension of their heightened interest in food.

So as we expect, people get really interested in food. Much less reasonable was the buying of old books, unnecessary secondhand clothes, knickknacks and other junk. Often, after making such purchases, which could be afforded only with sacrifice, the man would be puzzled as to why they had bought such more or less useless articles.

So that's essentially sort of saying people got really interested in certain purchases. Their preferences towards those kinds of purchases seems to have changed. They seem to have also some preference reversal. They seem to want to buy some things and then wonder afterwards why they did so, which all seems to be the consequence of being really, really hungry.

Then came the day when I lost my will to activity. I no longer cared to do anything that required energy, and days began to drag. That seems to be in some ways like a direct effect of like having a low caloric intake, but really, people's preferences towards physical or any other activity seems to have changed quite a bit due to hunger.

And another one here-- there's nothing that can hold my interest for long. I wait for meal times. And this is where people's attention seemed to really be lowered quite a bit, to the extent that you think attention affects people's preferences and decision making in various ways. Again, hunger seems to have this effect on people's choices.

Now, what's a more sort of like a real world example for you guys? Well, that's shopping on an empty stomach. There's sort of a classic study by Nisbett and Kanouse from 1969, but there's sort of other evidence later. And the folk wisdom is that shopping on an empty stomach leads people to buy more, and perhaps also more junk food.

Now, how do you study this? Well, you can think about different ways in which you might do that. You could sort of induce people directly to eat when they're hungry versus not, and there's various other ways in which you could think about like setting up experiments. Now, one very simple way in which one could study this issue is by just randomly giving a sample of individuals entering the supermarket a candy bar.

It could be a candy bar. It could be also a power bar or just some other food. And so now, essentially, what you do then is, you take a bunch of people who go shopping. And some of them might be hungry and some might not be. But then the ones that are hungry in particular, when they eat like a candy bar or any other food, they become less hungry. And so then we have a treatment condition that's less hungry and a control condition who is not getting any food who is more hungry.

And so then that allows essentially to look at variation in whether people are currently hungry while shopping. Another thing one could do is one could vary the timing of the last meal before shopping. There you worry a little bit, if that's not experimentally induced, that the people who go shopping right after having a meal versus not might be quite different in various other ways. But you could, in principle, also experimentally induce that.

And then you can monitor how much and what kind of food people buy. So hungry people tend to buy more. They just simply buy more food. They also more tellingly buy more junk food. So their preferences really seem to be different. To be clear, people tend to buy more, not just sort of like, say, I'm going to buy a Power Bar or, for example, a banana because I'm really hungry. I'm going to bite it right after I leave the store. But they buy significantly more, more than can be explained by just filling your immediate hunger.

And in particular, they also seem to change the type of food that they buy. In particular, in this case here, more junk food. Here is kind of what this then looks like. The left guy is full. He buys lots of lettuce, diet water, and so on. The guy on the right is hungry. He buys all sorts of things, a lot more as you can see, and a lot of also unhealthy food, including tortilla chips, cheese, sugar water, and so on.

So next time you go shopping, you should sort of introspect and see what your preferences towards what you buy have been affected. Now, another thing you can think about that affects your preferences or underlying state is your sleep. When you're tired, the whole world seems different in various ways.

One example you can think about is self-control. And sleep deprivation is associated with lack of self-control. There's sort of this old literature that says self-control is a muscle that replenishes overnight. There's some issues with this literature. And there's some experiments that might not necessarily replicate. But the idea that lack of self-control might be a consequence of lack of sleep seems very plausible in various ways. Notice that there could also be causality in the other direction. Lack of self-control might also be a cause of poor sleep.

Now, there's some evidence that sleep affects people's preferences. When people are sleep-deprived, they tend to gain weight. Tired people engage more in what's called cyberloafing. They surf more on the internet. There's also some evidence that lack of sleep might affect ethical behavior. People cheat more when they're tired. And we have some evidence in our experiment in India, which I'm going to show you in a little bit, at the end of the semester, that naps increase savings and seem to reduce people's present bias. So really this is some evidence of sleep affecting people's time preferences.

Another example is Badger et al. and addiction. This is a very small study with only 13 subjects. So you might take all of what's in there with a little bit of a grain of salt. But I think the underlying message of the findings are correct. So what did Badger et al. do? They elicited people's willingness to pay for a second dose of the heroin substitute, BUP. These are heroin addicts that are recovering. And they tend to have, often, cravings. And BUP essentially helps people not to relapse and not to take heroin again. And so, at regular intervals, they get BUP doses.

Now it will be sort of unethical to withhold people these doses. But what the experiment is doing-- and it's still ethically, perhaps, tricky-- but what they were asking for is people's willingness to pay for a second dose of BUP, the heroin substitute.

So all individuals in the experiments regularly receive their single dose of BUP. And then there's also an additional dose they could get. And then there was some variation in the experiment. What was the variation? There was a variation in the state of deprivation. So people were either asking all the questions about our willingness to pay for a second dose of BUP. Everybody always got the first dose. But then there was some variation in timing. For some people, the more deprived condition, people were asked about their willingness to pay for a second dose two hours before the scheduled first dose. And the less deprived people were asked right after the scheduled first dose.

Again, everybody's getting a first dose. And we're asking about willingness to pay for a second dose. Some people are asked about this two hours before the scheduled first dose, when they're really sort of deprived and craving, when they really want their first dose. And others are asked when they're less deprived, right after the scheduled first dose.

Notice that the second dose is always held constant. So everybody always gets the first dose. And the second dose is held constant in the future, at least in those kinds of questions. So really the willingness to pay should not depend on your current state. In the future, you're going to get a second dose, which is once you have already gotten the first dose. So whether you ask me like two hours before or right after your scheduled dose should have more affect on your willingness to pay in the future because that experience is held fixed. And you're always getting the first dose.

Now, they are also asked [INAUDIBLE], in addition, the timing of a potential second dose. So sometimes the second dose was later on the same day and sometimes the second dose was during the next week.

Now, what do Badger et al. find? They find that people's willingness to pay varies systematically with the state and the delay. What's the state here? The state is the state of deprivation, whether people are more deprived or less deprived. And then the timing is essentially either about today versus next week.

Now, the median willingness to pay for the second dose later today is \$50 in the satiated state and \$75 in the deprived state. And we're going to get back to this how we think about this, how we can explain this. But basically the way we explain this is that in the deprived state, people's willingness to pay is higher because they really, really want a first dose. So people's willingness to pay for the first dose is high in the deprived state and it's lower in the satiated state because they already have-- people's willingness to pay for any additional dose is high in the deprived state because they haven't even gotten the first dose. And it's lower in the satiated state.

But what people seem to be doing now, they seem to have a hard time imagining their utility in the future. That is to say, in the future, both people will be satiated. So what's happening here is that when people are in the deprived state, they really, really want BUP. And they're willing to pay more than when in the satiated state, again, for the second dose in the future.

Now, there's also some evidence of willingness to pay for the second dose next week. And again, people's willingness to pay-- and so the previous one was later today. Now it's for the next week. People's willingness to pay in the deprived state is \$60. In the satiated state, it's \$35.

So again, people's willingness to pay is higher in the deprived state. That is to say, when in the deprived state, really, really, really people want BUP. And they cannot really imagine that in the next week, when they're satiated, they will actually want BUP less.

To be clear, in both the first, when the dose is later today and when the dose is next week, it'll always be the case that people have already received the first dose. And then the question is always about what's your willingness to pay to receive a second dose. So when they're in the future, they will always be satiated. But people in the deprived state seem to be reacting or seem to be behaving as if, in the future, they would also be in the deprived state. But of course people will be in the satiated state. So people seem to overestimate, when they're in a deprived state, how much they would like the second dose in the satiated state in the future.

Now, here's another example by Schelling. And Schelling is a beautiful writer in economics. He's one of the best writers that economics has. Granted, there's lots of bad writing in economics. But if you want to read some beautiful writing by an economist, Schelling has a bunch of different books that are really nice to read and essays that are really beautiful.

Now in his paper in 1984, he talks briefly about a controversial question, which is the use of anesthesia during childbirth. And why is that controversial? Well, anesthesia is reducing the mother's pain during childbirth. But there's the potential for some side effects. And there's also the potential that people's experience is different. And for both those reasons, some people argue that anesthesia is not warranted and decide not to do that.

Now, in many cases, people's preferences change predictably over time. First, ex ante, before the woman gives birth, many women prefer not to use anesthesia. Once they are in excruciating pain, they request anesthesia from their doctor. And then, if they actually get the anesthesia from the doctor ex post, after the child is born, they regret their choice. They say, I wish I hadn't gotten anesthesia.

Now, that pattern holds even for women who have given birth before. And it leads to a bunch of ethical dilemmas and legal issues. For example, to which patient is a physician obligated, the one asking for anesthesia or the one who asked that it be withheld? The person told the doctor, I really, really do not want anesthesia. Even if I ask you later, please do not give me anesthesia.

But then, of course, the person that yells at the doctor and says, I really, really want anesthesia. Give me anesthesia right now. And so there's essentially preference reversals. People seem to want different things at different points in time. And the doctor can only make one of these two people happy. And so no matter what the doctor does, either if the person gives the person anesthesia, then person number 2 in my list here above is going to be happy, but person number 1 and person number 3 are unhappy. Or the person withholds anesthesia and then person number 2 is unhappy. So either way, the doctor cannot satisfy these different selves over time.

Now, can the physician enter a contract with a patient ex ante, before? Is it possible to have a legally binding contract here when he or she says, I'm not going to give any anesthesia, and you told me so? Do we want policies that make such contracts possible? Is that a desirable thing to do? Is it a legally required thing to do? Potentially impossible to require this legally. I'm not sure.

But the point here really is that there is some misprediction going on here. Ex ante, the person says, well, if I'm in pain, I will be fine. I'm not going to need anesthesia. But then, of course, when the person is actually in pain, the person wants anesthesia.

So essentially people seem to mispredict their preferences when they're in pain. Then that leads to all sorts of trouble here. I really recommend to read the paper by Schelling that's on the course website if you want to sort think about this a little bit more.

Now, so far, we talked about that preferences change predictably due to changes in underlying states. Now, one of those states, it would be like pain, hunger, sleep, and so on. Now, there's an additional stylized fact just that people make systematic mistakes at predicting preference changes. So not only is it that your preferences are different when you're tired versus not, but people are also systematically mispredicting that preference change.

Now, what is projection bias? To be clear, it's the fact that people underappreciate changes in their preferences, projecting their current preferences onto future preferences. So projection bias, to be clear, is not just some random prediction but a prediction with a systematic direction.

So people seem to understand the direction of preference change but not the magnitude. That is to say, when people are hungry and think about how it's going to feel when they're not hungry anymore, they kind of understand that eventually they will not be hungry and that maybe, suppose you're really hungry and you really want potato chips, you sort of understand that you don't really want that many potato chips when you're full. Nevertheless, when people are hungry, they buy 10 bags of potato chips as if they were hungry at least for like half of their life, even though they're going to be hungry for the next hour or two because then they're going to go home and go eat and so on. So people seem to understand the direction of people's preferences but not the magnitude.

Now the underappreciation of the effects of hunger on preferences is perhaps not the most economically important part of projection bias or the most important application. But there two reasons to consider this evidence. But one, it's perhaps the clearest evidence of projection bias. Second, people have lots of experience with changes in their levels of hunger. So any misprediction isn't due to the lack of opportunity to learn, really, when you should have probably learned this over time.

For some other instances-- this is only, for example, if you have lots of physical pain, or for childbirth, for example, you might say, well, that seems to be sometimes even "one in a lifetime" experience, where really it's sort of understandable that people mispredict their preferences when they are in excruciating pain. But that's understandable because you haven't experienced it before.

And even for somebody, a mother who has given birth to a child before, that mother might just misremember to some degree. But when it comes to hunger, people have been hungry so many times, and full, and so on. So really you had lots and lots of opportunities to learn what your preferences are when you're hungry versus not. So really it shouldn't be about a lack of opportunity to learn.

Now, people buy more on an empty stomach. That can be interpreted as a manifestation of projection bias. Hungry people act as if their future taste for food will reflect the current hunger at least to some degree, or more so than it actually does. But it's not completely clean evidence of projection bias, because there could be always other things going on. For example, if you're really hungry, some things might be more salient to you. And there you might buy certain things more because you pay more attention, for example, to potato chips and really sort of learn how exciting potato chips are and so on. But let's just go with hunger for now. There's lots of other evidence of projection bias in addition to the hunger example.

Now remember the paper that we discussed on food choices by Read and van Leeuwen. This is the paper where office workers are asked to choose between a healthy snack and an unhealthy snack. Now, so far, we had looked at this as a piece of evidence for present bias. We were looking at when people are choosing for the present right now versus when they choose for the future. Do people make more unhealthy choices when they choose for the present? That is to say the stylized fact that we found was that when people are choosing snacks for the future, they were quite likely to choose healthy snacks. You choose salad for the future as your snack in the afternoon. And then, you ask again, right now or today, what would you like, people seem to switch their preferences towards unhealthy snacks. They'd rather have, like, chocolates.

And so we interpreted that, at the time, as evidence of present bias. And surely it seems like that is evidence of present bias. Now, in addition, now, we're going to consider variation in the timing of those choices. So some people were asked when they were hungry, late in the afternoon. Others were asked when they were satiated or arguably satiated, immediately after lunch.

And then the snacks were to be received in one week. So that's held constant, the timing. Otherwise the timing is held constant. So we're always going to look at choices for snacks to be received in one week.

In some of the cases, the snacks were to be received when people were hungry or likely hungry, late in the afternoon. In some other cases, the snacks were received when people were satiated, immediately after lunch. OK, so there's like four cases here. There's the timing of choice. Either people are asked if they're hungry versus not. And then the timing of the receipt of their snack, which is when they were hungry versus not. And the question now is, are hungry people now good or bad at predicting their preferences when they're satiated, and are satiated people good or bad at predicting their preferences when they're hungry, in the different state?

So we're going to interpret here-- so here's a table from Read and van Leeuwen, table number 1. It's a bit of a sort of confusing table for no good reason. Let me just walk you very quickly through that table.

So the rows are the current hunger and the columns are future hunger. So the first row is when people were asked when they were hungry. So this is late in the afternoon. The second row is people were asked when they were satiated. That's like right after lunch.

And then the columns are-- again, remember, all of the choices are for next week. So the first column is choices for when people were hungry. That's late in the afternoon. And the second column is when people were satiated. That's right after lunch. Remember, again, all choices are for next week, so for one week in the future.

Now, we can interpret, now, the main diagonal of this table as reflecting people's true preferences. That is to say, here's no projection bias. People are asked, when hungry, to predict what they wanted or to say what they wanted in the future, when they were likely going to be hungry as well.

That is to say, 78% of people are choosing an unhealthy snack for the future when they're hungry-- so that's in the late afternoon-- for the future case when they're in the late afternoon. In contrast, only 26% of people choose the unhealthy snack when they're satiated-- that is right after lunch-- for when the snack will be delivered in one week, right after lunch.

So there's no projection bias here. Because essentially this underlying state, hunger, is held constant. Hungry people predicting for when they're hungry or asking what they want when they're hungry, 78% of people think they want unhealthy stuff. And again, satiated people, 26% of people say they want unhealthy stuff. So 74% of people say they want the healthy snack, in contrast, when they're satiated.

So again, that's what's written here. Late in the afternoon, when people are hungry, 78% of people choose the unhealthy snack for the late afternoon when they will be hungry. Immediately after lunch, when they're satiated, 26% of people choose the unhealthy snack for immediately after lunch when they will be satiated.



Now, then what we see here is the entries that are off the main diagonal. Those data fit the pattern of projection bias. Now what do I mean by that? Let's look first at people who are hungry but expect to be satiated. That is the first row here. Current hunger is hungry. That's essentially people are asked late in the afternoon. Or think about this as being likely hungry. And they're ask about why what do they want when they're satiated? That is right after lunch.

Now, remember, when people were satiated, 26% of people said they wanted an unhealthy snack. Now, when people are hungry, 42% of people say they want an unhealthy snack. So that fraction is higher.

So that is very much consistent with projection bias, the reason being that, when people are hungry and say what they want for themselves in the hungry state, well, 78% of people say they want the unhealthy snack. If people did not have any projection bias, they really should be saying 26% if you believe that satiated people know best what they want when they will be satiated. Instead, 42%, a higher fraction of people, say they want an unhealthy snack when predicting or when answering while they're hungry for when they would be satiated.

So people seem to understand the direction in which their tastes change as they become satiated, but they underestimate the magnitude of this change. Similarly, people who are currently satiated but expect to be hungry underestimate the effect of hunger on their preferences for unhealthy snacks.

Remember, hungry people, when they choose for next week when they will be hungry, 78% of people say they want the unhealthy snack. Now, in contrast, when they're satiated, 56% of people choose the unhealthy snack. So again, that fraction now deviates from the fraction when people are hungry, which is, again, consistent with people understand the direction in which tastes change as they become hungry, but they underestimate the magnitude of this change. So satiated people think that when they're in the future, hungry, they will behave at least a little bit like as if they were satiated. And hungry people seem to think that, when they're satiated in the future, they will still behave as if they were hungry.

Notice that this misprediction is partial. So again, both the 42% is in between the 26% and the 78% and the 56% is in between the 26% and the 78%. And again, the 56% is between the 26% and the 78%. So there's essentially some partial misprediction going on here.

Now, let me show you a number of other examples of projection bias. And then I'm going to write down a little model that sort of captures this information or the information that's here. So another example is catalog orders. This is Conlin et al. Imagine there's a cold evening. It's really cold outside. Maybe you were just outside and so on, and then you are online shopping. Notice that in Conlin et al., really, it's catalog orders and not online shopping because online shopping didn't quite exist at the time. But in any case, imagine you buy, online, a warm jacket. And again, it's really cold right now. You really wish you had a warm jacket. But then, when the warm jacket arrives, it's 30 degrees warmer when that happens.

Now, that could or could not be, potentially, a mistake. In particular, if it's an unusually cold day, if people have projection bias, they might sort of think that, well, right now they really need a warm jacket. And they can't even imagine that they might not need this warm jacket anymore in a few days from when it actually arrives from the order.

And so the hypothesis is that projection bias leads to an increase in purchases of cold-weather items on cold-weather days. And not only that, but controlling for receive-day temperature, the likelihood of returning a cold-weather object is higher when the day of ordering was cold.

So not only people are buying more cold-weather items on cold-weather days-- that seems very reasonable. Because you might just say, I always wanted a jacket anyway. It's really cold. And that just reminds me that I really need one. Or, for example, it's really rainy and you buy an umbrella. Because again, you just remembered that you really needed one or that you forgot it somewhere and so on.

The telling part here is that the likelihood of returning the cold-weather object is higher when the day of ordering was cold. And that's very much consistent with projection bias. People, on cold-weather days, act as if the future were only full of cold-weather days. And therefore their preferences would be such that they would always want to put weather items. Or put differently, people act as if their preferences for cold-weather items-- they really want a warm jacket-- would not change even if, in the future, there will be warm days.

Notice that, in some of this evidence, what's hard to distinguish is predicting people's preferences change in the future versus predicting future states. So if somebody predicted that it's going to be cold forever, who would act in the exact same way as if you predicted that your preferences are such that you always want warm jackets, you're always going to be cold.

Now, projection bias is about the latter. It's about people's preferences. But often it's actually hard to separate predictions of probabilities of how likely it is that it's going to be cold in the future from predictions of people's utility.

Another example is car purchases on Sunday versus radio or days imagine it's a nice, sunny day and you go car shopping as you do it. And the car dealer offers you an Audi TT. You know, "Why, I guess I never thought of the Audi TT." You go for a test drive. And the wind rips through your hair. It's really sunny and great. And having this car during a sunny day is really amazing. So you would love to buy a car like this. Why not?

Imagine, in contrast, car purchases on an icy day. Imagine you're car shopping around the time of a freak snowstorm. You kind of wanted to have a car, and somehow it's a snowstorm when you go car shopping.

Now, the car dealer says, how about this grand Jeep Grand Cherokee. And you might say, well, I guess I never thought of the Jeep Grand Cherokee. Why not try it? You go for a test drive and you gain traction on the black ice and you jump the curb with ease. It's just amazing.

Now, you might just say, well, I would love a car like this. Now, what people seem to be forgetting in both of these examples is that it might just happen to be a really bad snow day on that particular day. Or it might have just been really sunny on that particular day. And things like change over time. And people might not understand that their preferences for the Audi versus the Jeep might change with the weather over time. And that's exactly what Busse et al. are looking at. They're looking at how does weather impact automobile purchases. And in particular the classical prediction is that the weather, on the day of a car purchase, should have no influence on the type of car bought. So now what Busse et al. look at is do idiosyncratic weather conditions, controlling for time of year and predict car sales and then, in particular, also returns of those cars. And they look at two types of cars, in particular convertibles and then they look at four-wheel drives like the Jeep that I just showed you.

And the key part here is to look at idiosyncratic weather conditions. That is to say, controlling for time of the year. Suppose it's an unusually warm and unusually cold or icy day. Does that predict people's car sales or purchases and particularly also the returns of those cars? Now, the buying patterns are very much consistent with projection bias. People buy more convertibles on good-weather days and more four-wheel drives on bad-weather days.

And again, that evidence by itself is not necessarily projection bias. It could just be a good-weather day and you kind of always wanted to buy a really nice car, and therefore you just do it on a good-weather day. But people are also more likely to return their convertible if they bought it in the good-weather day. And they're likely to return their four-wheel drive if they bought it on a bad-weather day. And that's very much consistent with evidence of projection bias. On a good-weather day, you think your convertible is going to be amazing forever. You're going to love driving. And even if it's going to be raining, it's going to be amazing. But of course, if it's raining, that convertible is really not a car that you want to have. Or is really cold outside, that's really not a great car to have.

Similarly, if you have a four-wheel drive and it's really warm and nice outside, you'd probably much rather have a lighter car or potentially even a convertible. Now, one important question here is now, in any of the behavioral biases that we discussed is how does the market react to such behavior.

So one tricky part here is that you want to be kind of careful. So on one hand, you might say, well, let's exploit this kind of behavior. And I think that's right. You might get people to purchase stuff that they don't necessarily want. And surely that's happening in some cases. But to the extent that people can just return these items, that's actually not necessarily a great idea. And particularly if you're a car salesperson and you sell people stuff and then you just come back and you have to deal with them, and they're unhappy, and maybe, then, at the end of the day, you just return the car and actually don't buy anything eventually because they're unhappy, that's really not great for you. Plus it wastes a lot of time and effort.

On the other hand, sometimes people might just not return things. And they might just get stuck with their car and just leave it at that. And then it might be a good idea to exploit this behavior, particularly on a good-weather day, you might want to try to sell people convertibles. Another thing you could potentially exploit is essentially people just have higher willingness to pay. Suppose somebody really, really wants to buy a convertible anyway, regardless of the weather. But on a good-weather day, their willingness to pay might be higher. And so you might be able to sell them additional stuff for you might be able to sell them maybe a nicer version of the car, which probably they're not going to return. Because once they have it, they're just going to keep it. And so you might want to exploit that if you are a company.

There might also be some agency issues within the sales shop, which is like if the salesperson really wants to sell things and is incentivized to sell things, and that's regardless of the returns, then you might get into situations where, on sunny days, lots of convertibles are being sold but they're all being returned. So if you are a company, you kind of want to set the incentives right that the car salesman gets only rewarded for stuff that's not returned because it precisely gets you things that will not make you happy or profitable in the long run.

Now, there's another-- so as a customer, now, how might you take advantage of production bias? Well, there's an article here that you can look at that argues that winter is the best time to buy a convertible. And so the reason is, of course, convertibles are often cheaper in winter in part because the car dealer might have storage costs, the like. So the convertible is actually cheaper. Or just because people are taking advantage of in the summer when they really want convertibles.

Now, you might say, well, let's just buy a convertible in the winter. And I think that's, in principle, right if you have the time to wait until the summer or have a long time horizon. Now, you want to be a little bit careful with exploiting these market conditions as a classical agent, which is manifested by some of my colleagues who shall be unnamed who wanted to buy a convertible sometime in the summer. They got really excited about the convertible. But then they realized, well, winter is actually the time to buy a convertible because it's going to be cheaper.

So then they decided, well, let's just wait until the winter and then buy the convertible then. Of course, what then unfortunately happened-- or maybe fortunately. It depends on the perspective. Once winter arrived, they actually didn't want a convertible anymore. Because a convertible in the winter is really not a lot of fun. So projection bias perhaps also kicked in with my colleagues, where they then perhaps should have just bought the convertible and then predicted that, in the summer, they really wanted it. But instead they did buy it. Because wants a convertible in the winter?

So when exploiting behavioral biases from others, you want to be careful that you may be affected by it yourself. Let me give you one last example, which is this example by Van Boven and Loewenstein, which is about thirst. Here, visitors are asked before or after vigorous cardiovascular workouts-- so people who are going to the gym-- to complete a short survey. And the survey was as follows. Imagine that three vacationers in Colorado this past August embarked on a short 6-mile hike. As the day wore on, they realized that they were hopelessly lost. Worse, because they had packed lightly for a short hike, they had not carried much in the way of food or water.

And what people were asked, then, about-- and so they were essentially given this story. And then they were asked to do the following. In the space below, please take the perspective of one of the three hikers and describe your situation-- how you got into it, how you feel now, both physically and mentally, and what you are hoping will happen."

And now what happens and what the evidence of Van Boven and Loewenstein says is that thirsty subjects have way more empathy for others' thirst. So when you look at the different outcomes, before and after exercising-- remember, after exercising is when people are thirsty-- thirst was mentioned before hunger in the assay, for thirsty people much more. Thirst was unpleasant. For hikers, it was mentioned much more. Hikers would regret not packing water. Thirst more unpleasant for the self. Oneself would regret more not packing water. So all of these items that are sort of like thirst-related, people, when they were thirsty after exercising were way more focused on than not. So it is as if people have projection bias in the sense people understand much more that the condition of thirst is really bad when they're thirsty compared to when they are not.

OK, so let me take stock of what we discussed. So I showed you some evidence of projection bias for many short-run changes in preferences, and at least for some of them. There's hunger, thirst, pain, sleep, weather, addiction. There's other evidence as well-- arousal, anger, sadness, and so on-- when people aren't certain states, they have trouble to predict their preferences when they're not in that state.

Now, one key question that I don't have a great answer for you is why do people not learn? People have had lots of experience with a lot of these changes, in particular when you think about sleep or hunger or the like. But even for addiction, when people just have smoked a cigarette versus not, that happens all the time. So they really should have learned, over time, to predict their preference.

So the misprediction is really not due to a lack of opportunity to learn. Yet people really seem to believe, all the time, that this time is different, over and over again. So I think that's just a very deep cognitive bias in some ways in the sense of, when people are affected by certain visceral or other aspects that affect their preferences, it's really hard to imagine how they might feel when that visceral influence is not at play anymore in the future.

In addition to people underestimating short-run changes in preferences, people also underestimate adaptation to long-run changes. So Dan Gilbert, if you're interested, has very nice work, including a book called *Stumbling on Happiness*. And he gives many examples of the underestimation of adaptation, which they call immune neglect.

And so one example is how does a positive or negative tenure decision affect well-being? And so they ask current assistant professors at the University of Texas to forecast. And then they ask, as well, former University of Texas assistant professors to recall. And as a current assistant professor like myself, if you forecast, if there's a negative tenure decision, life will be terrible. And how are we ever going to live after being denied tenure?

And so when you look at what's actually happening, professors seem to be relatively accurate in predicting the immediate impact of the tenure decision, but they overestimated the long-run impact. That is to say, in fact, immediately after being denied tenure, people are maybe arguably or understandably disappointed and unhappy. But over time, the long run effect is much less severe. People seem to adjust to their circumstances. In the medium and long run, people are almost as happy as if not as happy as if they had gotten tenure. So there's actually not much of a long-run effect. Yet people seem to mispredict the effect. They really seem to think, in the short run, the effect will be really bad. And that bad effect will last forever. So people seem to essentially sort of understand there's a psychological immune system to bad events. People recover from negative shocks quite well over time. But people seem to mispredict this recovery, which which again is what's called immune neglect.

There's similar misprediction for other life events such as paraplegia or lottery wins, so for very negative and positive events. People seem to misunderstand that, again, in the short run, there's usually a pretty large decrease or increase in people's happiness and life satisfaction. But people tend to recover from that quite well over time. But in their prediction, people seem to mispredict that adjustment. People seem to think that the bad effects on happiness-- or bad or good effects and happiness-- would persist forever.

Let me now show you a simple model of projection bias by Loewenstein et al. that formalizes some of the intuitions that we discussed so far. So suppose true utility at time  $t$  depends on both consumption  $c_t$  at time  $t$  and the state  $s_t$  at the time  $t$ . So it's  $u$  of  $c_t s_t$ . The state could be anything that affects utility from consumption, ranging from the level of hunger or addiction, whether somebody had just smoked versus not or whether somebody has just gotten some BUP dose or not, or past consumption, et cetera, whether somebody is tired, and so on.

Now, the prediction at time  $t$  of future utility at the time  $\tau$  larger than  $t$  from assuming the  $c_\tau$  and state as  $s_\tau$  is  $\hat{u}$  of  $c_\tau s_\tau$ . And that equals  $1 - \alpha$  times  $u$  of  $c_\tau s_\tau$  plus  $\alpha$  times  $u$  of  $c_t s_t$ .

Now, what is this expression? Well, if you look at the first part,  $u$  of  $c$   $\tau$   $s$   $\tau$ , that is the correct future utility at time  $\tau$  from consuming  $c$   $\tau$  in state  $\tau$ . So if  $\alpha$  was 0, then this term here would go away, this term is just 1. So that is essentially a correct prediction. So there is no production bias. The person has rational expectation about their future utility at time  $\tau$  of utility  $c$   $\tau$  in state  $s$   $\tau$ .

Now, if  $\alpha$  is not zero, then there's this additional term here. So this is the weighted average of  $\alpha$  weights for this other term here, which we're going to talk about in a second. And  $1 - \alpha$ , the correct utility of consuming  $c$   $\tau$  in state  $s$   $\tau$ .

Now, what is this other term here? Well, the person gets  $c$   $\tau$  correctly, so the consumption is right but the state is wrong. And what state is the person using to predict their future utility? Well, it's using  $s_t$ , the current state. So the person is essentially using their current state to some degree to predict their future utility instead of the future state  $s$   $\tau$ .

Now, notice that if  $s_t$  and  $s$   $\tau$  is the same, then there's no problem here. There's going to be no misprediction. So it better be the case that the states  $s_t$  and  $s$   $\tau$  are different. For example, a hungry person would mispredict that utility from not being hungry. So the utility from not being hungry would be this part. The person is hungry right now at time  $t$  if  $s_t$  equals hungry. The person would predict the future utility even when they're not hungry, thinking that they will always be hungry in the future. So  $\alpha$  in between 0 and 1 is the degree of projection bias.

So just to give you more detail, the person predicts how she'd feel about consuming  $c$   $\tau$  in the future partially at least by how she'd feel about consuming it now. That's what the parameter  $\alpha$  is measuring.  $\alpha$  between 0 and 1-- it could be 0. It could be 1-- is the degree of projection bias. So  $\alpha$  equals 0 is correct understanding of future utility. That's essentially no projection bias. That's rational expectations.  $\alpha$  equals 1 is full projection bias. That's a person who thinks, essentially, their future state will always be their current state, which, of course, is not necessarily true.

So now that person will optimize according to a perceived future preferences. You had  $c$   $\tau$ ,  $s$   $\tau$ , and I'm going to assume for now is an exponential discounter. That's, of course, you could relax that easily. And, in fact, in some of the problem sets from previous years and from this year, that will be, in fact, relaxed.

Now let me give you a very simple example with hunger. Again, the problem set this year and past problem sets give you quite a few other additional examples that you can study.

So suppose there are two states. The state is either hungry or not hungry,  $H$  or  $N$ . And the consumption  $c_t$  is over burgers and money. So  $u$  of  $c_t$  where  $c_t$  is the number of burgers and money that the person has. And state  $H$  is 5 times the number of burgers that the person eats plus the remaining money that's left to spend on other things. The utility of  $c_t$  and when the person is not hungry is 1 times the number of burgers plus the money that's left.

Now what does that imply? Well, it implies that she's willing to pay \$5 for a burger when hungry and \$1 when full. How do we know that? Well, we know that because in the first case when the person is hungry, each burger gives 5 utils, and each dollar gets 1 util. So the exchange rate between burgers and utils is 5. So the person is going to pay \$5 per burger.

In the second case, the exchange trade between utility from burgers and utility from dollars is 1. So the person is willing to pay \$1 for a burger.

Now suppose alpha is  $\frac{3}{4}$ . Again, remember alpha is the degree of projection bias. I'm going to go back and show you. Alpha is this parameter here. Alpha is the degree of projection bias, so the degree of misprediction.

So suppose in our example here alpha equals  $\frac{3}{4}$ , and the person is not hungry right now. What is their willingness to pay for a burger tonight when she'll be hungry?

So if the person is not hungry right now, her correct utility or willingness to pay for a burger is \$1. We just established that already. She's full. And the correct utility for willingness to pay for a burger tonight when she'll be hungry is \$5.

Now her projection bias is  $\frac{3}{4}$ . So essentially she puts  $\frac{3}{4}$  of a weight on her current state and  $\frac{1}{4}$  on the full state. Well,  $\frac{3}{4}$  times 1 plus  $\frac{1}{4}$  times 5 or going essentially  $\frac{3}{4}$  between a difference of 4 is essentially 3. So the person essentially will think that their future willingness to pay is \$2 for a burger tonight. Or put differently, if she thinks she only wants like-- her willingness to pay will only be \$2 for a burger tonight, if she orders for a burger tonight and the burger costs like \$3 or \$4, she will not be willing to do that.

Now just to be very clear, so the person only essentially-- if alpha equals  $\frac{3}{4}$ , the person only incorporates  $\frac{1}{4}$  of the preference change from being full versus hungry in the future. So that is to say there's a difference of \$4 between \$1 and \$5, and she only takes into account  $\frac{1}{4}$  of that \$4, which is \$1. So willingness to pay is 1 plus 1 equals 2.

You can also write this down here instead of using the formula from the equation from before, which essentially is  $\frac{3}{4}$ , which is, again, alpha times utility using the current state, and  $\frac{1}{4}$  using the correct utility, which is the future state. Once you sort of calculate that, she would be willing to pay at most \$2 for a burger. So potentially she's making a mistake.

And so again in the problem set you will have more examples and questions that allows you to practice with this level model.

Now, why should you care about projection bias? I showed you pretty good evidence of projection bias in many settings in the sense of saying we think projection bias really exists and we think projection bias might be quite important. Sorry. It might be relevant in explaining people's behavior in those settings.

Now, the consequences in those settings is relatively unimportant. For example, in car purchases if rich people buy convertibles and then sort of return their convertibles, the loss from that is not particularly large. Maybe there's some redistribution across people. Similarly, the catalog orders, people might purchase too many things and return them. Really, the consequences are not that large, and the actions are often reversible.

So why should we then care about projection bias? Well, there's at least two really important applications that we're going to talk about for a little bit that really could matter quite a lot. One is addiction, which is people mispredict the utility from becoming addicted in the future. Second is depression and hope. When people are depressed, they might mispredict how they feel when they're not depressed anymore.

Some might argue that marriage and relationships are quite important or projection bias might be quite important. Particularly when people have fights, they might want to get divorced very quickly. They might sort of do rash things about getting married really quickly or try to get a divorce very quickly. And, in fact, there are some laws that often prevent that. People need to actually have some cool-down period for a few days where they have to sort of-- if one wants to get married, for example, you have to get a license, and it takes a few days, and only then you can get married. The same is true for a divorce, at least in some places.

Perhaps that's because people suffer from projection bias sometimes. And when they're really mad at somebody or something really bad happens between them or something really great happened between them, they might not be able to predict that their preferences and their views towards the other person might change over time.

But let me tell you now a little bit about addiction and depression and hope. So first addiction, and that's the main applications we're going to discuss.

So projection bias might be important for people's initiation when people are starting to take drugs. So here let's define the relevant state as the person's level of addiction. So how much has the person consumed of cigarettes in the last month or year?

In an unaddicted state, cigarettes are really not that hard to resist, right? If somebody has never smoked before, cigarettes are really not that attractive. Same is true also for alcohol, for example. In an unaddicted-- so in an addicted state, however, it's very hard to resist cigarettes. If you have smoked a lot recently, you really, really would like to smoke on any given day.

Now, the unaddicted person might sort of think that experimenting with cigarettes is fun but does not want to get addicted for the rest of her life. It's sort of fun to do it for a while, but really then you want to stop because being addicted for the rest of your life is really bad. You might get really bad health consequences, for instance.

And that might give you some false-- if you sort of have projection bias, you might have some false sense of control. If you project your current nonaddictive preferences into the future, you might think that you can stop smoking or taking drugs if necessary. And so then the person might try cigarettes, get addicted, and consume much of your life. Perhaps she wouldn't do so if she knew she could quit.

That's essentially to say if you're not addicted right now, you might predict that it's going to be always easy to stop smoking. And, in fact, right now it is really easy to stop smoking or resisting. But the person might mispredict that this will always be the case in the future even once she is addicted and once she has smoked a lot because for the nonaddicted person, it's really hard to imagine that it might be even harder to resist in the future.

In addition, there are quitting and restarting cycles that are very common when it comes to addiction. So addicts often express a desire to stop using substances permanently but are unable to follow through.

That's not surprising, and we've kind of thought about this before. And quasi hyperbolic discounting can essentially predict that as well.

Short-term abstinence in addiction is common while long-term abstinence is rare. So in 2000, for example, 41% of smokers stopped for at least one day trying to quit, but only 4.7% successfully abstained for more than three months.



That doesn't really seem like quasi hyperbolic discounting. So you wouldn't go through a pointless short-term pain and then not follow through if [INAUDIBLE] quasi hyperbolic discounting.

The reason being in particular that withdrawal symptoms tend to be, on average, strongest at the start of a quit attempt. So hyperbolic discounting would not sort of go over this really difficult pain in the short run if you don't follow through because it's essentially short-run pain.

And then three, recidivism rates are especially high when addicts are exposed to occasional cues related to past consumption. So treatment programs, in fact, advise recovering addicts to move to new locations and to avoid places when previous consumption took place. So let's sort of go through these one by one and try to see whether, particularly in two and three, whether we can explain those using projection bias.

So now let's define the state as the strength of cravings at the moment, and suppose this varies randomly or with exposure to cues. How do you explain the starting a quit attempt?

Well, suppose an addict is currently consuming regularly. She experiences periods of low cravings when it's easy to resist. That's like when the person has just smoked, cravings are low, and it's easy to resist. So then the person might think, well, it'll always be easy to resist, so she thinks it's worth trying to quit and starts the quitting attempt.

Now, how do we explain then how people abandon the quit attempt? Well, suppose the addict is currently on a quit attempt. Something triggers strong cravings, which might just be over time you get cravings or maybe the person might just be exposed to some cues. So then the person feels it's really hard to resist the drugs, and in particular, she thinks drugs will always be hard to resist because currently she's craving a lot, and you can't even imagine that it might be-- that that craving might stop over time. So she thinks the quit attempt is impossible to carry through, and so she abandons it. So that way projection bias can both explain why the person starts the quit attempt but also why the person does not follow through.

I'm sorry. Give me a second.

How do we think about the recidivism? Well, here there's not necessarily a projection bias needed. But here people might underestimate how important cues are in affecting people's utility. So if it's a case that an addict, for example-- when the addict walks by, like, a liquor store or when the addict is talking to friends who they have been taking drugs in the past and so on, that might really increase people's marginal utility of using drugs. So they might get really, really strong cravings, and it might be really, really hard for them to resist.

Now if you know that as an addict who is recovering, who's not taking any drugs currently, you should at all cost avoid these cues because you know you want to make sure that you're not exposed to them. But if you have protection bias, you might underestimate the importance of those cues in affecting your utility or your cravings, and so you might sort of walk into a bar thinking like, well, I haven't really been smoking or drinking for a long time, so I'm sure they will be able to resist.

But, of course, then once you're exposed to these cues, you will not be able to do so, and that miscalculation might really then cause recidivism. So again, in that sense, protection bias is also consistent with that kind of recidivism due to cues.

Now let me talk very briefly only about depression. And the reason being that we're going to talk about mental health in a future lecture.

So depressed individuals have the tendency to project the depressed feelings not only to the future but also to the past. And so in particular, the depressed tend to think about-- depression is the inability to construct a future. And in particular, when someone is depressed, the past and future are absorbed entirely by the present, and people can neither remember feeling better nor imagine that they will feel better in the future.

So that's to say if the person is really feeling terrible right now, they also think they always have been feeling terrible in the past, and they also think that they will always feel terrible in the future.

So then life might feel particularly hopeful if there's no scope for future improvements. And if you think your life was always that bad, there's really-- in some cases, people might think there's really not much of a reason to live.

And that's, of course, an extremely dangerous and difficult or consequential, potentially, misprediction because, of course, depressed people can get better through therapy and through drugs and so on. And often there's also remission over time. But if people mispredict that, they might, for instance, never seek treatment because they think, why would I ever seek treatment if I'm always feeling terrible for the rest of my life anyway? We're going to get back to that a little bit when talking about mental health.

Now, one thing that we have not really talked about is how to think about these different biases. In particular, what's the difference between projection bias and naive quasi-hyperbolic discounting? So both naivety and quasi-hyperbolic discounting and projection bias entail a misprediction.

So as an example, take a smoker who wants to quit. The naive quasi-hyperbolic discounter overestimates their future patience. So that person will sign up for a commitment contract to stop smoking and might fail due to the overestimation of usefulness of commitment devices. That person thinks that commitments-- the person understands that she has some-- suppose that's a partially naive person, which we need. Otherwise the person would never sign up for a commitment contract. Suppose the person is partially naive. The person will sign up for that commitment contract thinking that's going to help them stop smoking, but the person is going to underestimate how bad their self-control problems are. So then the person might fail because they think their commitment device is really, really helpful when, in fact, it's not strong enough.

Now for projection bias, the person might underestimate the influence of altered future states. So the person might also sign up for a commitment contract to stop smoking and then might feel guilty due to the underestimation of the changes in future cravings. That is to say if you are currently not craving, you might sign up for a commitment contract that would help you to continue not using drugs. And you might underestimate, however, how bad the cravings will be in the future once you see certain cues or the like.

So again, you might choose a commitment contract that is not strong enough because you mispredict your future preferences. In this case, not because it's in the future, not because you mispredicting beta like in case number one but rather because you mispredict how strongly future states will affect your preferences.

You can also think about another example, anesthesia during childbirth. You can get similar sort of patterns. Both naive quasi-hyperbolic discounting and projection bias might be able to explain people's choices there. In particular, both of them would potentially predict preference reversals.

Now, how can we tell projection bias and naivety regarding beta apart? Well, the key part here is that projection bias is a state-dependent misprediction. So people are more likely to predict future temptation to overeat when hungry. They're more likely to predict smoking when you haven't had a cigarette for a while. So when people are - when the state in the future is different from the current state, people might mispredict their preferences.

So now a present bias, in fact, has nothing to do with the state, essentially just about the future versus the present. So now what do we need? What kind of variation do we need to disentangle the two explanations? Well, you need variation in timing and in particular variation in states.

So present bias would say it doesn't really matter what state you're in. It doesn't matter whether you're hungry or not right now. There's always going to be a misprediction between the present and the future. And projection bias would very much say the states matter. It's a state-dependent misprediction.

So if you wanted to look at people's predictions for the future, present bias or naivete regarding present bias would say it doesn't matter whether you have smoked or whether you're hungry right now for your prediction for the future. Projection bias would say it does matter quite a bit.

But then a tricky question is when should you, in fact, offer people commitment devices? Well, you should probably offer people commitment devices if you want them to choose correct-- if you want people to make correct choices. You probably want to offer people commitment devices at times when they're, in fact, in the same state. So when somebody already has cravings, you want to offer them a commitment device because that person really understands how it will feel in the future.

If, instead, you offer somebody a commitment device when they don't have cravings, they might choose that commitment device, which I guess is good. But then they might fail because they mispredict how strong of a commitment device they might need in the future.

OK, so let me then briefly summarize. So what did we talk about? So we talked about state-dependent preferences. So preferences vary systematically with the underlying states-- for example, hunger. Food is tastier when you're hungry. Going on dates is less enjoyable while being sick. Classes are best when you're rested.

Now, people know that people's preferences vary with their states. They understand that. If you ask hungry people about are your preferences different when you're hungry versus not, most people would probably say they want different types of kinds of food and the like and they behave differently when they're hungry.

However, there are biases in state-dependent decision making. In particular, both intuition and psychology suggests that we fail to appreciate the extent to which our preferences change with the underlying states.

Projection bias is a specific psychological error of this type. People overestimate the extent to which future tastes resemble their current tastes, and they underestimate the influence that the state has on their utility, and that can lead to systematic misprediction. I showed you quite a few applications. In particular, addiction and depression might be the most consequential applications of projection bias.

Now there's one more bias that I told you I'm going to talk about very briefly-- again, that's going to be discussed in detail in recitation-- which is called attribution bias. Now attribution bias in some sense is very similar, but it's backward looking.

So what projection bias is the misprediction of the influence of future states. So when you think about the future, are you going to be-- how are you going to enjoy a meal depending on when you're hungry versus not or burgers versus salad? How is your future preferences shaped by future states?

Attribution bias is instead backward looking. It's the misprediction of the influence of past states. So that's to say-- and I'm going to sort of define that briefly-- when judging the value of a good, people are overly influenced by the state in which they previously consumed it.

Now let me give you some examples. And again, then you'll talk about this more in recitation.

People are more likely to return to a restaurant when first trying when hungry. So if you go to a restaurant when you're really, really hungry, you'll think the food is amazing while, in fact, it might just be that you're really, really hungry. So when you then come back to the restaurant, you might be surprised that, in fact, it's not that great.

People are also more likely to negatively rate a movie when they've seen it while tired. Well, why is that? Because the experience is just not that great, and then you sort of think the movie is really not that great. But, really, it's just you were tired.

People are also less likely to recommend a zoo to a friend if it rained during the last visit. Well, it's just the experience of that zoo is not that great. Of course, some days it rains and some days it doesn't. So you're not going to tell usually your friend to only go there when it's sunny. People are saying the zoo was just not that much fun.

And perhaps more relevant for you, people are more likely to-- or may be more likely to recommend a class that they took while well rested. So a few years ago I was teaching 1413 at 9:00 AM, and I got lots of complaints from students about the class being way too early. And I was telling them or trying to tell them that the class is really lots of fun. You may or may not be able to appreciate that if you really are tired.

So attribution bias might sort make you think that-- might attribute some negative or positive experience to the actual underlying quality of the issue or the good that's being offered as opposed to, as you should, to the underlying state. Again, you're going to talk more about this in recitation.

Now what's next? Remember, there is no class on Monday, April 20. On Wednesday, April 22, we're going to talk about gender, discrimination, and identity. Please read Heather Sarson's paper, section one, just the introduction.

Thank you very much.