

[SQUEAKING]

[RUSTLING]

[CLICKING]

**FRANK** All right. I'm going to get started. Welcome to lecture 16 of 14.13. We talked a lot about utility from beliefs and  
**SCHILLBACH:** how, in particular, anticipatory utility, utility about thinking about the future and what might happen in the future, may affect people's utility. And then how that in turn might affect how people, A, choose when to consume or when to engage in certain activities. Because anticipatory utility provides some motivation to push things forward, at least a little bit, so people can look forward to positive events.

Second, we talked about information acquisition about the idea that people, if they had derived utility from beliefs, that might affect how much information they might want to acquire. In particular, we talked about the idea that when you think about potentially negative information that you might receive-- in particular, we talked about Huntington's disease where people had a negative health information-- that they might not want that information with the motivation that they would like to make themselves feel better about themselves.

They might think that the healthier than they actually are, they might look forward to a healthy life at least for two more years. And that might depress their willingness to gather or receive information.

Now we talked a little bit about then a model and how to think about this. First, we talked about a model where people did not have the choice whether they could manipulate the information themselves. So we just talked about when somebody wants to seek or reject information, but we have the constraint. We had imposed a constraint. Let me go back for a second. We had a constraint imposed where we said, well, the person needs to be-- have-- hold correct beliefs conditional on the information that they have been exposed to so far.

So essentially, that was just an idea of, like, OK, here's some information that the person could gather. But the constraint was this number one that I'm showing you here is that beliefs need to be correct, as in like the person is a [INAUDIBLE] conditional on the information that they have received so far. And then we said, well, if you think about this, if this is just about wanting to gather information and be considered so far the case where the person could not actually affect future outcomes, so that's the case of Huntington's disease where there's no cure. There's nothing you can do about it. The question was just like, would you like to receive some information that's going to happen in the future or not?

Then we-- and I'm not going to go to this in detail very much. We then looked at the conditions that are involved for when does the person want to know about this. And the condition was essentially about this function of  $f$  of  $p$ , where  $f$  of  $p$  was essentially the utility derived from beliefs. And the condition was either if  $f$  is concave or  $f$  is convex. If  $f$  was concave, the person is information-averse. If  $f$  was convex, the person is information-loving and would like to find out about those things.

Now then we said, well, consider now also the possibility that people might be able to manipulate their beliefs. So you might say, what are the reasons? Why might you ever want to hold correct beliefs?

And in the above framework, if you could choose what your  $p$  is, there's really no reason whatsoever to not choose  $p$  equals 1. Remember,  $p$  here was the probability of your health being good, of being HD, Huntington's Disease negative, of not having the disease. So if there's no negative consequences of what's going to happen in the future, surely you have all the incentives in the world to deceive yourself and make yourself think that you are healthy, right?

And so now then the person was like, well-- and this is what happens [INAUDIBLE]. The expression that's here. There's essentially  $f$  of 1, and that's larger than  $f$  of  $p$  for any  $p$  that are there. And then if the future-- this is the term here-- what is actually going to happen in the future, if that's independent of your beliefs and the sense of that's going to happen anyway, you might as well make yourself believe that  $p$  is large.  $P$  is 1. You're healthy.

Now then what we left things at was the question of, why might you not want to choose  $f$  of one anyway? Why might you not want to say  $p$  equals 1? And that's a question to you.

**AUDIENCE:** If your belief of what's going to happen might affect the future outcome in a negative way, you-- I guess if there were a treatment available for Huntington's disease, you would rather know whether you have it.

**FRANK SCHILLBACH:** Right. So Huntington's is not the greatest of all examples that I chose to start with, because there's no cure. But you're right. For example, in particular, for diseases such as HIV or the like, surely it would be very helpful to know. Because then you can engage in proper treatment that might actually help you get better. Presumably, you only get the treatment if you actually know what the disease is like.

In the specific Huntington's disease, we also talked about a few things. We talked about other actions that you might be able to take. For example, if you are able to save for retirement, if you are able to go to travels before you get sick. Or there was questions about whether people want to have children, questions about like what partners they want to be with and so on and so forth.

So if there's a bunch of other economic choices, or other choices in life that depend on whether the person is positive or negative, it seems like that person should want to know. And then deluding yourself might get in the way of making optimal decisions. But broadly speaking, just to summarize, if there is no action item, as in like if there's some information about the future that where really, even if you knew the information, there's nothing you could do about what's going to happen in the future, it's not obvious why you not want to just delude yourself and think like things are rosier. The world is looking better than it actually is. Because it might just make you happier, at least in the moment, even if things in the future might turn out to be bad.

And so this is what we already just discussed. So incorrect beliefs can lead to mistaken decisions. Well, that's correct, but sort of overly positive beliefs are an economically important indication of utility from beliefs.

That is to say on the one hand, you might-- and this is what we just discussed-- you want to believe that you're healthy if it makes you feel better about yourself or the currently or the future. So you might want to convince yourself that that's the case. But this is what we just said. Overoptimism distorts decision making in some ways-- for example, health behavior, whether you want to seek treatment, but also whether you want to adjust to potentially bad events and some other ways in economic choices. So you might not want to-- so you might not want to be overoptimistic because of those distortions.

And then the optimal expectations will then trade off these two things. On the one hand, you're healthy and maybe happier right now, thinking you're healthy. On the other hand-- or thinking that the future will be bright anyway. On the other hand, it might make-- change your choices in some bad way.

Now, there was another reason that people mentioned, which was potential for disappointment, right? If you think the future is always going to be great, even if you can't effect it in any way well, at some point reality is going to set in, and then you might get really disappointed. We haven't really talked about this very much. That's a way in which you can think about perhaps people potentially choosing their reference point.

So if you think people have referenced dependent utility, and if you're really overoptimistic very much in a way that allows you to potentially choose your expectations and your reference point. And then you might not want to be overly optimistic, because your reference point might be then too high and any outcomes that you're gonna eventually receive you might sort of view worse because you expect and believe things to happen. So to the extent that being true, or being overly optimistic leads to disappointment, you might not want to engage in overoptimism, even if there's nothing you can do about the outcomes, because of that potential disappointment.

But in general for decision making with anticipatory utility, at least some overoptimism leads to higher utility than realism. Because essentially it makes you feel better in the moment. But there's this trade off potentially with how good you feel in the moment versus how disappointed you might be in the future in case you're too overoptimistic. Any questions about this?

Oh, sorry. OK, so then let me show you some other overoptimistic beliefs in addition to the evidence [INAUDIBLE] showed you. So there's a classic study by Weinstein [INAUDIBLE] that asks students to make judgments of their students' chances for a number of outcomes. It's similar to the question that I had asked you at the beginning of class.

There are two measures that Weinstein uses. One is what's called comparative judgment. This is-- excuse me. This is how much more or less likely the average student thinks the event will happen to them relative to the average student.

And the second one is what's called the optimistic/pessimistic ratio, the number of students who think their chances are better than the average classmates divided by the number who think their chances are worse. So these are both sort of measures of overoptimism. They measure slightly different things, but they're broadly-- they're fairly highly correlated.

So what does Weinstein find? Clear evidence of overoptimism when you look at different things about stuff that's going to happen in your life, owning a house, salary of larger than \$10,000 a year-- I think this is 1980. That's been a while ago-- traveling to Europe, living past age 80 and on. The comparative judgment is always fairly high.

So that number would be 0 if people were, on average, realistic. Remember, comparative judgment is how much more or less likely the average student thinks the event will happen to them relative to the average actual average student. And the optimistic/pessimistic ratio, that should be one if people were not overoptimistic. It's also clearly a large N1.

Remember, that's the number of students who think the chances are better than the average classmates divided by the number who think their chances are worse. But it works for positive events, but also works for negative events. So when you ask people about drinking problems, suicide, divorce, heart attacks, all sorts of bad things, people think they're now, of course, less likely to happen to them, these kinds of events.

And then they end up optimistic/pessimistic ratio, that's positive. That's just because it's like the optimist flips, so people are way more optimistic, or-- about-- they think they're not going to have drinking problems, get divorced, get heart attacks, and so on. So people tend to be very optimistic about their future lives compared with when you asked them to compare it to others.

Now there's other examples of that. Couples believe there's a small chance that their marriage will end. Small business owners think their business is far more likely to succeed than their typical similar business. Smokers understand the health risks of smoking but don't believe this risk applies specifically to them.

There's a long list of kind of these kinds of behaviors. People tend to have very-- hold rosy beliefs about their future, even in the presence of a sort of objective information. So even if you give people information about the health risks of smoking, they understand the health risks. But then they're like, well, that's not going to apply to me.

Of course, there's no good reason to actually dismiss. It should apply to anybody. So people just want some things to be true presumably because it makes them happier in some ways.

So then beyond future prospects, people also tend to have overly positive views about their abilities and traits. So, for example, 99% of drivers think they're better than the average driver. 94% of professors at the University of Nebraska think they're better teachers than the average Professor at the University. I don't know what this number looks like at MIT, but probably it's larger than 50%, as well.

So [INAUDIBLE], you can find this kind of evidence from a range of domains. In some cases, there's some other potential explanation for this, but broadly speaking there's a pretty robust evidence that people are overconfident, overly positive about, A, what's going to happen to them in the future, and B, about thinking about their own skills and abilities.

Now if you look at MIT students, however, in contrast when you ask about students different questions, and this is from a recent survey from a few years ago that MIT asks students. One question here is academically, I would consider myself above average at MIT. People say about 31% of students say they're above average. In particular, female students tend to think they're below average.

People might have biased beliefs about the average. So what you think is an average MIT student might actually be like the fifth or whatever, some of the highest percentile, because essentially you're exposed to all the success and math Olympics and whatever, wins that people receive, which is just not what the average student is like. There's, of course, also very severe selection in terms of there's so many smart and brilliant people at MIT that, in a way, people's perception is quite biased.

The questions here were asked about the average at MIT, but even there I think people might just be biased in terms of thinking what's the average.

There might also be some worries about disappointment in the sense-- and this is what I was saying before-- when you think about being overconfident, maybe in some ways-- so one issue with, or one reason why you might not end up being overconfident for a long time is if you receive feedback. And at MIT, people receive lots of exams and so on where you can learn or interact with others, but you can notice how people are doing overall. And if you're worried about being disappointed, or you have been disappointed several times already, then thinking you are the smartest person in class might not be a good idea because you might get disappointed again.

And then we'll talk a little bit more about gender in a different lecture later. That's a very common theme-- one second-- is that female MIT students, or in general women, but in particular also female MIT students, are particularly under-confident. We don't exactly know why that is, but it's a very robust finding, not just at MIT. In some ways, then, there's still a question-- so it's a very nice observation that in a way, beliefs can also be a form of self-motivation.

So if you are really interested in academic success and being really good at exams and so on and so forth, if you are under-confident, that can be a motivator in saying that you're really worried about failing. You're worried about doing badly and so on. Then you work extremely hard, and then you do better in exams. You surprise yourself positively. That can be an important motivator.

Of course, that could also go the other way. It could be really discouraging. If you think you're really terrible at everything, then you might-- at the end of the day, you might not study at all anymore, because what's the point?

So there's a bit of a trade-off here as well. And it could be that if you think that people are-- when they're under-confident, that leads to motivation and it makes them better at school. Plus, they're going to be less disappointed and get positive surprises. That could be an explanation.

I think there's a bit of a question again. Why is that the case at MIT and not in the overall population? But it might have to do with frequency of feedback, and school, and so on and so forth. And maybe also just with social norms or other things that maybe are somewhat trickier to explain.

I was teaching Harvard undergrads as a grad student. I'm not so sure that Harvard students are under-confident, to be honest. I haven't seen the survey evidence on that, but my sense is that MIT students are more under-confident compared to Harvard students. Let's leave it at that.

But I think it's true if you look at Caltech and other type of schools that are very similar to MIT. I think-- my guess is you'll find similar results in those kinds of surveys.

OK. So the next thing we're talking about is very quickly is what's called ego utility. So partly I want to talk about anticipatory utility, which was the idea that people want to look forward to good things in life. And therefore, they're-- have inflated beliefs about what's going to happen in the future.

A very related aspect is what's called ego utility, which is-- so just to recap, inflated beliefs might be explained by anticipatory utility essentially to say higher ability means better future prospects, and people want to convince themselves that they have high ability. And therefore, if you think you're really smart, that means good things are going to happen in the future. But it might also just be, like, right now, you're just feeling better about yourself, and therefore you have inflated beliefs about yourself.

So that's often tricky to distinguish empirically in the sense of like, I might just feel like I'm good looking, smart, and so on and so forth, because it makes me feel good right now. Or I might want to hold those beliefs, because that means good things is going to happen in the future, and I will have a bright future going forward. So these things are slightly tricky to disentangle, and both of them probably play an important role at the end of the day.

So now one thing I'm going to talk about a little bit more also later in the semester is about mental health. And so there is this literature that argues that positive illusions are-- promote, in fact, psychological well-being. So that's the argument that is here, that overconfidence is, in fact, vital and important for maintaining mental health.

So overconfidence and that makes people happier. People are better able to care for others if they think they, themselves, are doing better. They're better at doing creative and productive work. That's also better over confidence helps people manage negative feedback.

So overall, all of those things are [INAUDIBLE] so this literature argues that overconfidence actually helps people with their everyday lives. In particular, helps people maintain a positive-- a good mental health. And there's this term called depressive realism, which is to say that realistic expectation might actually be detrimental to mental health. So that the world is just difficult in various ways and bad things are happening in many lives, and if people are not overly optimistic about what's going to happen in the future and so on, that might be actually bad for their mental health.

We'll return to this issue on happiness and mental health in a later lecture, but I wanted to flag as these things. They're very much linked to each other. And another reason perhaps why overoptimism might be, in fact, good or not detrimental is because not only makes people happier, but it might also help people protect their mental health.

OK. So then just very briefly, and this is more like-- there's some research in some of those things. Some of this is more anecdotal. What are some factors that affects the extent of positive biases? So people tend to have greater biases, in fact, with respect to prospects and traits that are personally important to them.

On the one hand, you might say, well, people should also be better informed about those things. But to the extent that it affects their utility, they might be more biased. If it's stuff that's really important to you, you might be more biased because that's precisely what's going to make you happier.

Available or imminent objective information tends to decrease biases as in like-- and this is what I was talking about with MIT students. When you think about academic performance, you get lots and lots of feedback every single semester. So in a way, it's more difficult to be overly optimistic, or to maintain an overly optimistic picture of your academic self if you've got all these objective signals over and over again.

So one prediction here would be that, for example, if you looked at first, second, third, and fourth year students, that perhaps some of the under-confidence only comes later in the semester. Of course, there's other reasons why that might not be the case.

Then if feedback about the prospect or any prospect is more ambiguous and subject to interpretation, biases tend to be greater. So in particular, people are particularly biased in situations where people get ambiguous feedback where you can interpret things either way. So either you were really smart and did something really great, or you're just lucky.

And people tend to, in many things in life, people tend to interpret sort of lucky things that happened in their lives, tend to think that those are due to skills. And there's this very nice work in a book by Robert Frank-- he's at Cornell University-- who sort of argues that people tend to interpret a lot of things that happen in their lives that really are just mostly luck. They tend to think that it's due to their amazing abilities and so on.

And that's the reason for that, often, is that a lot of this feedback, or a lot of this information is ambiguous. You can interpret it either way. And people like to think that good things are because they did really great, and bad things are because they just happen to be unlucky.

If people-- and that's quite interesting. If people feel like they have control over the outcomes, biases tend to be greater. Expertise sometimes increases biases, but not for experts who get very good feedback. So meteorologists or the like who get lots of feedback all the time, they actually tend to not be overconfident, because, again, if you get so much feedback all the time, it's hard to sort of maintain your overoptimism.

OK. So now one question you might have is, well, I showed you a bunch of data information on biased beliefs but not a lot on actual action. Of course, the testing behavior, and the Huntington's disease, and so on [INAUDIBLE] some actions. But a lot of the stuff on the Weinstein and other studies were just about self-appointed beliefs.

And one question you might have is, well, this is about self reports. What about your preference on choices that people make? So there's some evidence that's from lab experiments, and there's some other evidence that I'll show you that is about actual choices.

So one very nice experiment is the paper by Eil and Rao from 2010. The way this works is people are given feedback about people doing IQ tests, and they're rated according to their physical attractiveness by others in the study. And they get feedback about this IQ test score and their physical attractiveness.

These two things are chosen presumably because people care a lot about them. People care about how smart they are. People also care about how good looking they are.

And then so in the study, the authors elicit people's prior beliefs, the beliefs at the beginning before they receive information, about a rank between one and 10. So people are put in groups of 10 and then asked about, in a group of 10 people, how do you think you rank compared to others?

And then people get some information, and they get, in particular, bilateral comparisons with another participant in the group. And then the authors elicit their posterior, so beliefs after receiving this information, and also the willingness to pay for true ranks. And broadly speaking, what the authors find is there's asymmetric processing of objective information about the self.

In particular, when people receive positive news, so when somebody says you've been compared to another person in the study and you're better looking, or you have a better IQ score than this other person, people tend to be roughly Bayesian, as in they tend to be pretty good at actually updating based on that information. In contrast, when people receive unfavorable news, they tend to essentially discount the signals, and they tend to essentially not really react to that information.

So that's very much consistent with people who are very happy to receive positive feedback and actually good at integrating it, as in they're good at doing the math as a Bayesian would want them to. But once they get negative feedback, they essentially mostly just discount that information, presumably and very much consistently with motivated beliefs. People tend to want to have positive images of themselves.

There's also some evidence of people's willingness to pay for information if they think the information will be good, and not willing to pay for information if they think the information will be bad, so if they're low in the ranks. Again, people want to hear good things about themselves, and they learn more or they're willing to learn more when they receive positive news. They're also willingness-- willing to pay more for information that they think will be positive.

There's some other evidence that's similar to that. There's a very nice-- or a similar paper by Mobius et al that's very similar to Eil and Rao. There's some very nice work by Florian Zimmermann in motivated memory. So that paper argues that it's not just about when people update immediately. So when people are given positive or negative information, they update differently.

But what Florian Zimmermann shows is that in his work that when people are given this information, in fact, in his experiment there is no asymmetric updating in the short run. So when people are given this information about-- this is about how they did in a test in Germany, in the short run people seem to actually be not overly optimistic in their updating. But then a month later when the author goes back to people and asks them about the information that they have received and their beliefs, people tend to remember positive signals more than negative signals.

And that suggests that memory can play an important role in the formation of motivated beliefs. When stuff happens, good stuff, people remember more, and bad things, not so much.

Now those are all lab type experiments where people are incentivized to give correct answers. But in a way, it's somewhat contrived. There's also quite a bit of evidence in the real world setting, in particular, when it comes to finance or trading behavior.

For example, there's Barber and Odean that shows that-- looks at overconfident small scale investors. Men, for example, tend to trade a lot more than women. We also know from other settings that men tend to be more overconfident than women, and those men, then, also lose more money or make less money plausibly due to overconfidence.

With this type of evidence, it's often a little tricky because it's hard to say, is it due to gender? Is it due to overconfidence and maybe other things that are correlated with this? But there's quite a bit of evidence on that kind of behavior that essentially overconfidence leads to worse decision here in this case trading decisions.



There's also a very nice paper by Ulrike Malmendier and Tate on managerial hubris. They have a clever way of identifying overconfident managers, which essentially is managers that are overconfident-- they are holding stock options, if they own stocks. So you're not supposed to do that, and you will most likely do that only if you think your company is going to do really great compared to other companies. Instead, you exercise-- you should exercise your stock options and diversify and not to be overly invested in your own firm. But if you think that you're really great or your firm is really great, you [INAUDIBLE].

And precisely those overconfident managers engage their business in more mergers, and those mergers, Malmendier and Tate show, in fact, are not good. People tend to do too many mergers that essentially destroy value, presumably due to overconfidence. So there's that kind of evidence as well.

So overall, we think there's quite a bit of evidence of people systematically holding overoptimistic beliefs. And those overoptimistic beliefs are affecting people's choices in a way that makes them worse off. In some ways, in particular, in finance choices where people tend to lose or destroy value.

OK. Any questions on this before I move on to heuristics and biases? OK.

So then we talked about three of our four issues here already. The last thing that we haven't talked about very much is to say Bayesian learning is just really hard, and people might just not be good at it, not for any motivated reasons or because of any utility reasons [INAUDIBLE] anticipatory [INAUDIBLE] utility. But it just might be really hard to do, and people essentially use heuristics because it's too hard to compute these things on their own.

Now we talked about this already quite a bit. Lots of economic choices are made under uncertainty. Almost anything in life involves some uncertainty, so we need to know something about the likelihood of relevant events.

And if you decide which topic in the course to focus on when you study for an exam, and you only have like one night to do so. If a basketball coach decides whether to leave tired players in the game, you kind make some probabilistic judgment about which player is going to be best. Many medical, managerial, educational, career decisions, essentially all those kinds of decisions are involving probabilities based on probability of good things happening, some probabilities of bad things happen, and you need to make some estimates about those probabilities.

Now how do make individuals probabilistic judgement of this kind? So far when we talked about risk preferences earlier in the semester, we essentially looked at, how do people make these choices for given probabilities, right? That was essentially prospect theory. That was about Kahneman and Tversky, 1979, which was how do people think about risk?

And what we always assumed was that the probability distribution was given. There's a 50% chance of something good happening. There's a 50% chance of something bad happening, and how do people make those kinds of choices? What does their utility function look like?

Now we're going to talk about, how do people learn about such probabilities? How, in the first place, do they update what probabilities are like? Which is, of course, a very important input in their decision making.

Now, this is based on pioneering work by Kahneman and Tversky, even earlier than the '79 paper. This is, in particular, a paper on 1974, which is in the reading list and on the course website.

Now first, I want to start with-- and this is important in behavioral economics to acknowledge and try to emphasize people are pretty rational and pretty good in various ways, in the sense of people get lots of things right. In particular, people are pretty good in terms of broad directions of people's probability judgments.

Here's a very simple example. Imagine you're deciding whether to see the new James Cameron movie. This could be a James Cameron movie. It could be a Korean drama or anything else.

People think about they want to watch something new. Well, how are you going to do that? Well, you can read some online reviews. You hear some opinions from your friends. You look at some Rotten Tomatoes or other ratings, and then you try to make those kinds of choices.

And so you start-- maybe you probably would, like, OK, have I seen a James Cameron movie before, or I said Korean drama, whatever you're interested in. have I seen something like this before? How much did I like the previous one?

That could be your prior that you start with. And then you get these other pieces of information of online reviews, various opinions from friends, or a Rotten Tomato ratings, and how do you use those, then? What do you do with those?

So essentially what you're going to do is you take some-- you got some prior, which is about James Cameron movies or whatever to start with. And that's like that-- that's what you use to start with in the first place. And then you get essentially these pieces of information.

And if they're positive pieces of information about those specific movies, then you're going to update up. And if they're bad, if your friends don't like it, as you said, they update that.

Notice that it could also be that you have a friend who has terrible taste in movies, and if that friend does not like a movie, that you update the other way and so on. So it doesn't need to be that your friend's movie ratings are positively correlated with yours. That could still be informative. If somebody has terrible movie taste and likes something, that could be actually good news for the particular movie.

But exactly-- this is what I have here. Essentially what you're going to do is a version of this would be you start with your base rate or your prior. And then essentially you use the various pieces of information and then adjust your probability up or down. Suppose you have to sort of [INAUDIBLE] willingness to pay or your probability of seeing the movie or doing other things, or your probability of liking the movie, you adjust it up or down based on the different pieces of information that we get.

And so generally people are pretty-- and this is a pretty reasonable and mathematical, mathematically well-founded procedure. This is actually pretty close to what a Bayesian would do. So people are actually pretty good at this. They're pretty-- they understand the basics of forming likelihood estimates. They know the direction in which features a baseline likelihood and available information should affect estimates. So overall, people are pretty good.

Now what's much harder to do is not just a direction, but also the extent to which you should update. That is to say, people know which direction to move, as in like if you and I like similar movies and I tell you this movie was great, it's pretty clear that you should move upwards. But what's much harder to understand is by how much should you move upwards in your estimate of how good the movie or the probability that you like the movie is.

And that's true for many things in life. People are pretty good at the direction in terms of where to move up or down. What they're much worse at is how far to adjust their estimates. And that's where they should be using Bayes' Rule, but that's extremely cognitively demanding. And so in many situations people tend to not get that right.

Now, just to sort of remind you, and I think in recitation you discussed this as well, what is Bayes' Rule or how does Bayes' Rule work? Let me go very quickly through that and then focus on that next evidence.

Here's one example. Suppose you have a coin that you start off thinking as a-- is fair with probability of  $2/3$ . That means the coin is biased towards heads or the coin is heads and tails. It's biased towards heads with probability  $1/3$ , in which case heads comes up 75% of the time.

Now if you flip the coin and it comes up H, heads, what's the probability of that, that it's fair? Clearly, it should be  $2/3$ , because you've got some signal, but by how much should you adjust is actually very hard to do. So you can think about this in your head, and you realize it's actually kind of hard to do, unless you can actually write it down.

So suppose you had to make a very quick decision. In particular, it's actually very hard to do this quickly on your own. You can do this particularly with one signal, but suppose you get several signals and it gets much harder and harder over time.

Now how would you do this? You can sort of-- in this particular case, you can do a graphical illustration. Again, I'll do this very quickly, but you can read about it overall.

You can think about, there's sort of two types of coins or like urns. There's a fair coin that's represented by an urn with equal number of heads and tail balls. And there's an unfair coin, which is represented by an urn in which 75% of the balls are H. So you don't know whether the coin is fair, so you don't know which urn you're drawing from.

And one way to think about this is imagine you're drawing from an urn that contains both of those urns. Of course, you need to then also take into account the probability of the fair ball being  $2/3$ . So the fair urn has twice as many balls as the unfair one. So you can sort of represent it in that way.

And then suppose you draw a ball randomly and it's H, what's the probability that it came from the fair urn? Then again, if you've taken probability classes and so on, that should be relatively simple for you to calculate. But here in the graphical way, you can look this very simple.

There's essentially seven red balls. There's four in the fair coin and three in the unfair coin. So the probability that it came from the fair coin is  $4/7$ . Similarly, if you draw a T, then the probability it came from the fair urn is  $4/5$ , right?

And that's, in some sense-- you can write down the math, and it's pretty clear and so on. But in a way, it's hard to do this in your head in the sense that you know it should be lower than  $2/3$ , but how much lower and how much should you adjust is actually kind of hard.

And here the answer happens to be  $4/7$ . And so here's sort of the formal way of doing this. I'm not going to go into this. But trust me, the answer is  $4/7$ .

Now, even this simple example of Bayes' rule, maybe it was very simple for you. But it's somewhat difficult to follow in the sense of would you have gotten the first sentence right immediately if I gave you 10 seconds? Maybe, maybe not. But probably quite a few of you would have gotten it wrong.

Now, in addition, then, if you try to apply these kinds of rules or the base rule in new situations with multiple pieces of different information, it's far more difficult. Suppose I told you to draw like 17 times and get those signals and without replacement and with replacement and so on. Things would get way, way more difficult.

And the urn example is, in fact, a much easier one. Because there are at least I'm telling you exactly what the probabilities are. But in many real-world situations, you don't even know that. A friend telling you the movie is good or bad, it's not clear how to read that signal, how do you think about that signal as a whole.

So but most people don't know the precise rule of Bayes' rule. I guess, at MIT, most students do. But in the general population, people don't. And even if they do, they can't or don't want to think so hard in many cases. So what people do instead, they use intuitive shortcuts to make judgments of likelihoods.

Now, that's a good thing. So we want people to use these shortcuts because otherwise they will just freeze and not be able to make any choices. And so, in some sense, it's good that people are sort of simplifying problems and they make some choices quickly. So in some sense, it's good. But it also leads to systematic mistakes. Because in some sense, you can use shortcuts. And they help you to get things approximately right. But now, if we can understand their shortcuts, then we can also understand, potentially, their systematic mistakes.

So what people tend to do is they focus on one or a small set of the situations at hand that seems most relevant. They focus on that to make their decision. And often, then, they systematically neglect other more complicated other issues of the situation, which makes the likelihood of their estimates typically incorrect. And now the question is, can we sort of understand what people focus on? And can we understand what causes these systematic biases to sort of make clear predictions on how people do things wrongly. And then, perhaps, if you want to help people make better choices, you can provide them with valid information as well.

Now, one particularly interesting issue is sequences of random outcomes. So as I just showed you, if you sort of drop balls from an urn, that's a very relevant situation, not because people in the real world in real-world situations, drop balls from urns, but rather because, in many situations, you do actually get repeated signals over time that you should use to update.

So an investor might observe past performance of mutual funds before deciding which one to invest in. A patient of a doctor might observe the outcome of prior surgeries before deciding whether to undertake the surgery or which doctor to choose. A coach can observe the recent performance of a basketball player before deciding whether to put the player in the game. There's lots of kinds of situations where people get repeated signals over time and then have to try to infer likelihoods of probabilities of certain events to happen. So that's a very common thing in the world that we should try and understand.

Now, one systematic pattern that we see in the world is what's called the gambler's fallacy. That's the false belief that, in a sequence of independent draws from a distribution, an outcome that hasn't occurred for a while is more likely to come up on the next block.

So the important part and what's doing a lot of work here is independent draws. Suppose there's a distribution of outcomes and there's independent draws, which means that, essentially, what happened in the last two, three, or four draws should not have any effect on what's going to happen in the next draw.

So if you play roulette or any sort of poker or the like, if you get read several times in a row, that has exactly no impact on how likely it is that red or black is coming up in the next draw. But people tend to have sort of this almost like folk knowledge of, well, if red came up a few times, now black is due, and the other way around.

And that's true in many different situations. But the important part here is that there's a sequence of independent draws. So as to say, these are situations where the probability distribution is actually null. That is to say, when you play roulette or you go to a casino, it's not like you need to learn about what's the probability that the roulette is fair or not. You know that there's some monitoring and so on in those situations. You know, essentially, that the chance of getting red or black is the same. And yet people tend to make these kinds of updates.

There's a nice paper by Gold and Hester, an old paper that shows this quite nicely, in which subjects are told the coin with a black and a red side would be flipped 25 times. And the experiment is slightly sneaky here, because essentially actually they reported people a predetermined sequence of events. So there's some deception involved here, which is a little bit trickier but it's mostly fine.

So what they see is essentially the subject sees 17 mixed coin flips. So they see a sequence of coin flips. It's red and black and so on. And then, at the end of it, you see essentially one black and four reds. And so, again, this is supposed to be like independent draws, which to make this case, in this experiment, they sort of rigged the last five draws, which is not quite ideal. But let's just go with this for a bit.

And so what essentially people see in this experiment is mostly evidence of red and black where roughly red and black is even. If anything-- so the 17 are mixed-- these are actually draws-- if anything, people saw more red than they saw black. Then, on the 23rd flip, after seeing these 22 flips to start with, the participants were given a choice between 70 points for sure or 100 points if the next flip was their color. And points are essentially stuff they can get money for eventually.

And so now it's randomly chosen whether half of the subjects color was red and half of them was black. And so now the propensity to take the 70 points reveals the beliefs about the odds that the next flip would be their color.

So if you think the probability is 50%, which is kind of like, given the information that you got, you should probably roughly think it's 50%, and if you're risk-averse in particular, you should truly choose the 70 points for sure. Because an expectation, if you get  $p$  times 100, if it's 50%, you get, essentially, 50.

So you're only got to choose the "100 points if the next flip is your color" if you really think your color is likely to occur. And in particular if your probability of that occurring is like 70% or higher, notice that that even assumes risk neutrality. So if people are risk-neutral, if they think the probability is 70%, they should be different. And again, if they're risk neutral and they choose number 2 here, their color, well, that's only the case if they think the probability of the color occurring is higher than 70%.

And so given the evidence that they saw, it's going to be rarely the case that-- or people should really, given the evidence that's there, the best thing they can do is essentially-- or the thing is set up such that they really should choose number one. But of course, the idea is here, if people think that it was four times red, so now black is due, the prediction would be that people would be more likely or will choose item number 2 if their color is black.

And so 24 of 29 red subjects chose to take the sure thing, which essentially is option 1 here that I showed you. And eight of 30 of the black subjects did. So these are essentially the people who choose their color. So if people have black as their color, they essentially choose this option because they really think now black is due because there's been essentially four reds in a row.

Now, again, if you're told-- and that's a little bit in play in this experiment-- about these being independent events, which for coin flips, really they should be independent events, that doesn't make any sense. You should, if anything, think red is more likely because the coin might be sort more likely to be red. But sort of thinking that black is more likely surely is not a good idea. And it might essentially make you lose money.

Now, there's some interesting variant of this for which they set, for some variants, the 23rd coin flip was delayed by some time, 24 minutes. And what they then find is, in fact, weaker evidence of what's called the gambler's fallacy, which seems to say that once you let the coin rest for a while, people seem to think that letting the coin rest for a while, which sort of makes it more likely, I guess, that the streak continues, or less likely that essentially black now is the case.

So people seem to these fairly-- and I think that's fair to say-- irrational beliefs about or biased beliefs about what's going to happen next. And that seems to be easily swayed by even relatively small things. It's like the coin needs to revert. But if you wait for 24 minutes, not so much anymore. There's quite a bit of work on the gambler's fallacy overall in various settings. That's just one setting to illustrate this. There's quite a bit of other work in that area. And it's a fairly robust finding that people have found in the literature.

Now, a second pattern that people find is what's called the hot-hand fallacy. That's the idea that, in particular basketball fans or other sort of types of sports fans-- fans, players, coaches-- believe that there is systematic day-to-day operations in players' shooting performance. And that's for basketball, but it's also true for other sports. So the idea is that the performance of a player may sometimes be predictably better than expected on the basis of the player's overall record. And so what people would say is, well, the player is on fire today or is a streak shooter and so on.

And so the idea is that "on fire" today means that he or she is more likely to hit his or her shots than on other days. So that's to say prediction is that made shots should cluster together. Like on one day you happen to be really good, maybe, in the first half. The second half, and conditional of having made a few shots, the next shot should be more likely than sort of the unconditional probability.

There's lots of work on this issue, starting by Gilovich and Vallone and Tversky in 1985. People have gone back and forth between saying in fact there is sort of such a thing as a hot hand or not. The initial claims were essentially there is no such thing. People believe that there's a hot hand going on and players are really running hot, but in fact, they're not. Some other evidence later showed maybe there is actually such a thing. And it was contradicted again.

So it's kind of complicated. But overall it seems to be that people tend to be-- players, fans, and so on-- tend to be quite overoptimistic about the hot-handed streaks happening in reality.

Now, one question is that that seems kind of odd in some ways. And so, on the one hand, the hot-hand fallacy and the gambler's fallacy seem to be opposites of each other. That is to say, the gambler's fallacy is the belief that the next outcome is likely to be different from the previous ones. If I show you, if you're gambling, you've got like four reds, now you say it's going to be black next. So if you have three heads, then you think that your coin is going to show some tails, that's essentially saying you saw a bunch of outcomes of one kind, the next one will likely be different.

On the other hand, then hot-hand fallacy is saying, well, it's a belief that the next outcome is likely to be similar to the previous ones. Now, what's going on here, you can think of both of these as a consequence of what's called the belief in the law of small numbers. Now, what is that? You should all know the law of large numbers, which is, in large samples of independent draws from a distribution, the proportion of different outcomes closely reflects the underlying probability. So essentially once you have sufficiently many draws from some distribution, the distribution will converge to what you're drawing from.

Now, what's the belief in the law of small numbers? Well, it's the belief that, in small samples, the proportion of different outcomes should reflect the underlying probabilities.

So usually, the law of large numbers, it's called the law of large numbers because you need the large number of draws. That's why it's called that way. But people seem to think that the law of large numbers also applies to small samples. And so how does it explain our puzzle? Well, then there's a question of does the person know the underlying distribution?

So if you think you know the underlying distribution-- for example, if you think for sure that the coin is fair or the roulette table is not cheating you, you might say, well, roughly, I should see as many reds as I should see blacks or I should see as many heads as I should see tails. And so now, if I have a sample of like four draws, if I already have like four heads, then I might sort of think, in that small sample, I need to see, essentially, a roughly equal fraction of heads and tails. And therefore, if I have seen four reds, the next one is due to be black. And of course that's not true because essentially the large number does not apply to small samples because it's called the law of large numbers and not the law of small numbers. But mistakenly, people seem to think the law of small numbers applies.

In contrast, if the person does not know the distribution, the belief in the law of small numbers can lead to the hot-hand fallacy. That's to say, if you're trying to learn about a distribution of shots or the underlying probability, if you see some good events in a row, if you see three heads or tails and so on, if you see a player making three shots in a row, you might sort of try to infer, oh, today is a good day or a bad day. And people tend to over-infer. They tend to think they can learn more from these three shots than they actually do. Really, to be able to estimate whether a player is running hot on a given day, you need quite a few draws. But people think, if a player makes three shots, that already means that his probability of making good shots on that day are way higher than it actually is.

So essentially people mistakenly seem to over-infer, from a small number of observations, what the underlying probability distribution is. And in that sense, the belief in the law of small numbers can lead to the hot-hand fallacy as well. So these are essentially two very different conclusions that seem to come from the same underlying issue.

Then the last pattern I'm going to show you-- by the way, there's more patterns such as those, but I'm sort of just showing you three of them to give you a sense of what's going on. Remember the example that I showed you in the first lecture, which was a question about base rates, base-rate neglect, which is the question about 1 in 100 people have HIV. If we have a test that is 99% accurate, if a person tests positive, what's the probability of having the disease?

And so I showed you this before. The true answer is 50%. But quite a few if you answered 99%. And the plausible explanation is that people probably forgot to take into account how few people in the population have the disease to start with.

And so this is what's called the base-rate neglect. When people are given some new information, people tend to-- in this case, I guess, a test being 99% accurate, people tend to focus on that piece of information. They tend to neglect the underlying base rate, which, in this case, is like one in 100 people tend to have the disease to start with. And again, that's a natural consequence of focusing too much on one central aspect offhand, which is the 99% accuracy of the test. In a way, that, again, can be useful. But it's sort of systematically getting people to make wrong choices. And there's quite a bit of evidence showing this in the literature.

And here's, again, what you should be doing and here's the probability of this. Trust me that this really is 50%.

Now, the last piece here, in terms of biases or systematic biases, what I showed you so far was, well, people have a hard time making right choices even in simple informational environments. Now, two-dimensional learning is even more complicated. And what I mean by two-dimensional is to say, if you try to learn about news or the accuracy of news, if you watch news online or anywhere else, you don't need to only learn about somebody gives you some information and you try to update based on that information, but rather you also need to learn at the same time about the accuracy of the piece of information that's given to you. That is to say, if you watch something online, any article, not only you need to sort of understand what to learn from that information, that piece of information, but you also need to understand the underlying source

Of it. And there's quite a bit of work recently on fake news and the question on, can people in fact detect fake news? Does it have to do with partisanship in particular? So you might think that Republicans or Democrats one certainly news to be true. So you might sort of think that Republicans are more likely to want to believe views that are biased towards Republicans, Democrats are more likely to believe news that are more likely to be Democrats. And therefore, they might be more likely to trust fake news or fall for fake news.

David Rand, in fact, seems like it has actually-- and at least in their work seems to have-- by the way, he's at Sloan, in marketing, in the marketing group-- he seems to find that it's less about partisanship but rather about people being not paying attention and being somewhat lazy and making their choices. And once you draw their attention to, here's an article that might be fake, people might seem to be actually quite good at learning about the accuracy of it.



Now let me just summarize for a bit and then see whether there's any questions. So what I've shown you is essentially using Bayes' rule is difficult, or I gave you some sense of it. I showed you three systematic deviations from Bayesian updating, which is the gambler's fallacy, the hot-hand fallacy, and base rate neglect. These are all well known deviations from Bayesian learning.

Now, 1 and 2, I have argued, can be explained by the belief in the law of small numbers. And there are a really important real-world implications of those types of outcomes. For example, one very nice recent paper-- Maddie talked about it a little bit in recitation, but only very briefly-- is this is paper by Kelly Shue and co-authors about decision making by judges, umpires, and loan officers, where essentially, if you're a judge or a loan officer, you see a bunch of applications over time. And then judges seem to essentially engage in what's the gambler's fallacy. That's to say, if they see three applicants in a row that happen to be very good or very bad, that tends to affect the following applicant.

That is to say, judges, when they decide about parole and so on, if they have given parole several times in a row, they might think the next person who is in line, if they have granted parole several times in a row, the next person then would be less likely to get parole, even though of course these are independent events. It just happens to be that one person happens to be after or before a certain applicant in a random way.

So there's really important decisions that the gambler's fallacy might affect. And more recently, when you think about base-rate neglect, in some sense, the example that I showed you seems to be a fairly academic example in some ways. But when you think about-- and I don't want to talk too much about COVID-19 since you see that in your life too much already anyway-- but if you think about tests for COVID-19, there is now talk about antibody tests. And particularly as tests that seem to be highly informative.

In some sense, their sensitivity and specificity, meaning the type 2 errors are pretty lower. Their sensitivity and specificity is very high. These tests tend to be highly accurate. But in fact, how much you can learn from those tests is actually very limited if the overall fraction of people who are infected is not that high.

And that's exactly the example that I just showed you. If the fraction of people who are positive in the overall population is relatively low, you are doing a test. And essentially, receiving a negative signal actually is not giving you a lot of information overall. And people tend to miss that part precisely because of base-rate neglect.

Let me stop here for a second. That was a little bit fast. So I want to see if there's any comments or questions. And then, only in the last 10 minutes, talk a little bit about heuristics and biases. And this is what Kahneman and Tversky really are very well known for.

So the way to think about, then, biases and probability judgments is what we already discussed. The starting point is that applying the laws of probability and statistics is often impossibly hard. So people use their quick and intuitive judgments to make likelihood estimates. And so there's a seminal work by Kahneman and Tversky and lots of subsequent work that tries to think about these biases. If you're interested in learning more about this, *Thinking Fast and Slow* by Danny Kahneman is really a terrific book about thinking about this.

Now, what's a heuristic? It's an informal algorithm that generates an approximate answer to a problem, quickly. And therefore, because it's informal, it's kind of hard to model.

Well, the previous things that I showed you before-- let me just go back for a second-- the gambler's fallacy, hot-hand fallacy, and base-rate neglect-- here, you can write down models that capture these phenomena pretty well. For example, for the base-rate neglect, you could just have a parameter like how much weight people put in the base rate, how much weight do people put in the new information that they get. And then that parameter will essentially-- you can estimate that parameter and you can write down a model that's not Bayesian but close to Bayesian.

Similarly, for the hot-hand fallacy and gambler's, you can model that people essentially apply the law of small numbers. And again, that's a model you can write down and estimate. In contrast, some of the stuff that I'm showing you next is much more difficult to model because in that sense, some basic laws of probability do just not apply anymore in the sense that people don't respect them anymore when they make their belief updating.

Let me show you, in a second, what I mean by that. But again, heuristics have a good and bad side. So they speed up and make possible cognition. So they help you make decisions in your everyday lives. And without heuristics, it would be really hard to make any decisions overall.

But because they're shortcuts, they occasionally produce incorrect answers or biases. So these are essentially unintended side effects of adaptive processes. And that makes it very useful to study heuristics and biases together. In the same way as you would study vision and optical illusion together, studying them jointly is very helpful because they're sort of the product or the result of the same thing. Things are hard. Therefore, they use heuristics. And because they use heuristics, that leads to systematic biases.

OK, so now I'm going to show a few of those heuristics just to give you a sense of what these are. One of the most well-known ones is what's called the representativeness heuristic, which is about Linda. Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice and has also participated in anti-nuclear demonstrations.

The question becomes, then, rank the following statements from most to least probable. And so here's a statement. You can read about them, all sorts of different things. In particular, what Kahneman and Tversky are focusing on is items 6 and 8. And item 6 is "Linda is a bank teller" and item 8 is "Linda is a bank teller and is active in the feminist movement."

Now, one thing that should be clear is that 6 is more likely to an 8, because essentially 8 is sort of conditioning on two statements, while 6 is only conditioning on one. So if Linda is a number 8, of course it's implied that, if she's a bank teller and something else, of course then she also is a bank teller.

So should always be the case that 6 is more likely than 8. But what people tend to do quite consistently is that, in fact, people rank it more likely that Linda is both a bank teller and a feminist than that she's a bank teller. And there's quite a bit of work on this. And showed this in various ways. But essentially it's what's it the violation of the conjunction law, which is a very basic law of probability. And precisely because it's such a basic law of probability, it's hard to actually model this because essentially you can't use simple things that should be true in probability theory.

There's some potential concerns about this, people misunderstanding this. But there's a subsequent experiment that essentially shows that that's not the case.

Now, what's going on here is what Kahneman and Tversky called the representativeness heuristic, which people essentially use similarity or representativeness as a proxy for probabilistic thinking. So based on the available information, people form a mental image of what Linda might be like and then ask about how likely it is that-- for example, she's a schoolteacher-- they might ask them how similar is my picture of Linda to that of the school teacher?

And so that turns this similarity judgment into a probability judgment. So here, of course, the example is very much rigged in the sense of they were asking things about discrimination, social justice, and anti-nuclear demonstration and so on, which made people think, well, that's a person that's likely to be part of the feminist movement. And then people focus on that when making these choices. And they forget the fact that, essentially, 6 must be, by definition, more likely than 8.

So that's a common thing that people do. So there's quite a bit of evidence of that.

And again, that's a very reasonable heuristic. And it probably works in many cases. But it also leads to very predictably bad choices. And it's a poor predictor of true probability in several situations.

Similarly, what's called the availability heuristic, people assess the probability of an event by ease of which instances or occurrences can be brought to mind. So when you ask people, for example, about are there more suicides or homicides in the US each year, people will think about suicides and homicides that they can recall. They might think about what they watched on TV and so on. And they judge the frequency of each based on how many instances they can recall. And people will be much more likely to recall homicides because they are much more salient in the world. And that leads people to think that murders are way more common, which in fact is not the case.

Again, it's a very sensible heuristic that people use. More often than not, it's easier to recall things that are more common or probable. So that's reasonable overall. But again, there are, predictably, things that are then more salient, that get more attention, people tend to think they're much more likely than they in fact are.

Let me skip the familiarity and anchoring and adjustment and just sort of like conclude. So what have you learned about beliefs overall? So we studied several reasons why people might miss information and fail to learn we talk about attention. Attention is limited, and therefore they might miss things.

You talked about, why might they miss important things? Well, because they might have wrong theories of the world. Then we discussed people might derive utility from wrong beliefs and therefore might actually not want to learn, or might be systematically engaged in trying to deceive themselves. And then we talked about people who might be bad at Bayesian learning.

Now, what's important here, I want you to take away it's important to understand these underlying reasons why people are misinformed because that might lead to vastly different policy implications. For example, if you were to make information salient, make sense, if you think people miss information and important things in the world-- so if you, for example, think that people just might not be aware of the fact that certain foods have more calories than others, then making that information very salient to people when they purchase things, drawing their attention to that stuff, could be really important and could be really powerful in helping people make better choices.

Or if they have wrong theories of the world, they think calories are not that important or if they miss the fact that smoking causes cancer, well, then we should really sort of draw their attention to them and help them understand it. And that's what a lot of labeling and a lot of making things salient often is about.

Now, but, if people don't want to learn, if people actually have motivated beliefs in certain ways, and in fact they know that, in some ways, deep down inside, they know the relevant information already-- smokers, for example, might know exactly that smoking is bad for the health. So providing them with information about smoking will not actually make them learn.

Or like if you wanted to give somebody information about how to best take care of your health when it comes to Huntington's, well that's not helpful if the person doesn't actually want to believe that they have Huntington's. They're never going to read what you send them if they want to maintain their positive image about their positive beliefs about their health.

And in fact, when people derive actually from beliefs, correcting those beliefs can make them worse off. If somebody wants to believe that for the next-- they know the probability of having Huntington's is reasonably high, but they want to maintain the beliefs that they're healthy and they want to be happy for the next 10 or 20 years, and then maybe, at some point, the disease will break out. But their choice is they want to think that they're healthy and they do not want to adjust their behavior.

So who are we then to tell them otherwise that's their choice? If it makes them less happy, we might actually make them worse off. So in some ways, there's at least some argument for respecting people's choices if they want to be deluded, if they want to delude themselves, and the consequences are not severe in the sense of these situations where they can't really do much about that, then in fact leaving people uninformed might be the right thing to do.

And then understanding systematic biases based on kind of help improve decisions. That's, in some sense, a lot less controversial if you think that people are systematically making wrong choices because they learned wrongly and they just have wrong information in their heads, sort of providing them with the correct information seems like unambiguously something [INAUDIBLE].

That's all I have on beliefs for you. Next time, we're going to talk about projection and attribution bias. But I'm also happy to answer any questions in the next few minutes that you have on this lecture or on the summary that I just showed you