

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: Welcome to Lecture 15 of 14.13. Today we're going to talk about utility from beliefs. Overall, lectures 15 and 16 are talking about utility from beliefs and learning.

Well in the previous lecture, we talked about attention and the idea that attention might be a limited role-- so you might not have the capacity to attend as much as we can. And then thought a little bit about, what kinds of things are people paying attention to? And can it be that people systematically attend to the wrong things and miss important things in the world?

Today we talk about different deviations from, perhaps, optimal belief or optimal information acquisition, which is people might derive utility directly from beliefs that therefore have potential incentives to deceive themselves in certain ways. Next time, we're going to talk more about learning-- sort of deviations or systematic deviations from Bayesian learning.

So to summarize today, we're going to talk about the utility from beliefs. People directly deriving utility from beliefs, from what they think about the world, about themselves, what's going to happen in the future or how smart they are and so on. On Monday, we're going to talk about non-Bayesian learning-- this idea that people essentially have trouble being visual learners.

Being Bayesian is quite hard. You may have taken probability theory and so on and have learned quite a bit. But even as a very smart and educated MIT student, there will be Bayesian learning problems that are way too hard once you have several variables, lots of information.

It's just really hard to do these things in your head. And people might just not be very good at it and then sort of use heuristics and biases instead.

And then after that, we're going to talk about projection and attribution bias, which essentially, if people have trouble predicting how they might feel in different states of the world-- if you're hungry, it might be hard for you to not think about being hungry or think about how it might feel. But you're not hungry. If you're sad, it might be very difficult for you to understand how it feels when you're happy and so on.

And so I should say, on Thursday and Friday, tomorrow, you're going to talk in recitation about Bayesian learning, which is sort of like, how do economists or statisticians think you should optimally learn? What's the Bayesian benchmark here? And we're going to talk then about-- on Monday-- deviations from that.

OK. So the first big picture overview is like, why might people miss information and fail to learn? What are potential reasons? Why might people do that? And we talked about true potential issues already last time. One is, very broadly, attention is limited. That's to say, there's so much information in the world that we just cannot attend to everything. Right?

And so one specific example was inattention to taxes. But you could think about there's lots of different prices in the world. It's just way too much information in the world for you to attend to everything.

And so one example that we talked about was inattention to taxes. It's saying that taxes that are not included in the sales price might be easily missed by people. And we have this very nice paper by Chetty et al. that demonstrated that-- that people essentially systematically under-estimate or under-appreciate taxes that they are exposed to.

People pay more and more attention to taxes that are directly included in the sales price compared to taxes that are only added at the counter. And of course, both of them you have to pay. And so you really should not really under-appreciate those taxes, because that's just expensive for you. We did not get much into why are people doing that, but more like show the existence of people seem to be missing those taxes.

Then we talked about reasons why people might systematically-- and even in the presence of lots of data in front of them-- not update properly, even if they are, in fact, Bayesians. And the reason might be that they might have the wrong theories of the world.

You might think that certain variables are just not important for your well-being or for any important things in the world, and therefore, you have essentially-- since your attention is limited, you will only collect some information. You only focus your scarce attention and memory to what you think is important things in the world.

And so now if you do that, you might systematically miss or not collect data that would help you essentially improve your theories and therefore, even in the very long run, even when you get lots of data, you will not update. Notice that that's quite different from rational inattention.

There is rational inattention theories would say, well people's attention is limited, but they focus attention to whatever is most important for them. So it can't really be that they have huge losses from a lapse of attention, because if they were, then they would direct their attention to something that could be potentially important.

In contrast, the "Learning by Noticing" paper, the theory paper by Hanna et al. that we discussed, provides a reason why people might not pay attention to stuff that's really important for them. And so, again, if your theory says that certain aspects or certain information is not helpful, you would not pay attention to those pieces of information. And if you don't pay attention to those pieces of information, you will never update your own theories and your own theory will essentially persist for a long time, potentially forever.

We're going to talk now about two other reasons for potentially wrong beliefs. One is anticipatory utility. This is, essentially, people like to look forward to good things in the world. Suppose you have a vacation coming up in six months from now. You might think about this positive event going forward. Essentially, good or bad events in the future that might happen and you might derive utility from those events already now.

Second, there's what people call ego utility, which is, people derive utility from thinking that they are smart, good-looking, and so on. So essentially, they derive utility from thinking positively about themselves.

Now for both of those kinds of pieces of utility, people now potentially have some incentives to delude themselves and to think that things are perhaps better than they actually are, because that will make them happier. And so then often there is, at the end of the day, there's some trade-off between having overly positive or optimistic beliefs will make people happier.

They feel good about themselves. They feel about good things happening in the future. But that may come at the cost of not preparing or taking wrong actions because people are sort of overly optimistic. Right?

So one example would be people might not get tested for certain diseases because they want to think that they're the healthy people and want to be happy about the future and look forward to a positive future. But of course, if they were getting tested, that would help them to take optimal actions-- potentially medication or any other actions reacting to the potential disease that they may have.

And for ego utility, for example, if you always think that you are smarter than you actually are, well that might make you happy. But at the same time, that might have come at some cost and you might not prepare properly for some exams or interviews or other things. You might just miss important things that might hurt you along the way.

And then finally, the last reason that we'll talk about next week is people might simply just be bad at Bayesian learning. Then there's no utility from beliefs or the like involved. It's just a lot of the computational issues that we ask people to do are just really hard and people might be bad at that. And therefore, people use heuristics and biases.

And then systematically, most of the time these heuristics and biases are quite helpful and help people making reasonable decisions. But then often these heuristics and biases are, in some ways, systematically wrong. In some cases-- and we can then think about some systematic errors and perhaps offer some ways in which you can improve people's decision-making by avoiding these errors.

Now let's talk about utility from beliefs. So it's important to understand that economists typically define utility functions over outcomes, such as money, consumption, health. Think of these things as things that people consume one way or the other, or things that they can usually measure at some point.

So you can look at it like, you have money. Then if you have money, you will spend it on consumption goods. It could be apples and bananas, of course. It could be also like haircuts or other services or health, for example.

And I can measure whether you have good or bad health. Usually we think about some outcomes that could be at levels of outcomes. So essentially, like three apples or five apples or seven apples are the arguments of utility function.

It could be also-- as we discussed when we talked about reference-dependent utility-- it could be in part being deviations from some reference point or changes over time. But importantly, here, the arguments-- or goods and services, essentially, are the arguments of the utility function.

Instead, another source of utility could be beliefs about such outcomes. So now it's about what might happen either in the future or what's happening right now. And you might directly derive utility from such beliefs. Now such utility from beliefs can be a very powerful source of utility.

One example, for example, would be a high-profile public speech. Suppose you have to give a public speech in front of the entire universe. Now of course you might sort of derive utility directly from that experience itself. You might like it a lot. You might find it very stressful.

So the 10, 20, 30, whatever minutes-- however long the speech is-- might give you positive or negative utility. But perhaps the utility of that experience itself might be or could be dwarfed by the stress and anticipation derived beforehand from anticipating it.

That is to say, you might think about it. You might prepare about it. You might worry about it. You might not sleep at night.

So there might be lots of stress or positive excitement and anticipation. Every day-- if you give the speech sometime in September, every day until September you're going to think about it positively or negatively. And you might derive positive or negative anticipatory utility from it.

There's another part of utility that we're not going to talk about. It's utility from memories. So you might afterwards look back and say, oh, that was really nice. And you have a video, and you're going to enjoy it, and so on. We're going to not talk about that, but that's sort of-- backward-looking utility could be quite important as well.

Now when you think about utility from beliefs, one way to think is that this is all in your head. But it turns out the utility from beliefs can actually even affect physical outcomes. In particular, there's a large body of research on placebo effects, which is defined as a treatment that can help patients merely because he or she believes it will. That is to say, if you give people a placebo or a sugar pill, something that essentially has no active ingredient, and tell them that this is, in fact, an effective drug for something, that drug itself, relative to control, might have treatment effects.

Now that often is the case for things like pain medication. Although usually we think pain medication-- if you give people placebo pills, compared to actual pain medication, these drugs are about a third as effective as pain medication. Depends, of course, on the dosage. So that's quite powerful, but then you might say, well that's only in people's head, and maybe people just say they're happier or less in pain, when, in fact, they're not.

There's also some studies that have actually found some physical symptoms. So you give people a placebo pill and there's actually physically, in some measurable ways-- hard measures, not just self reports-- people are doing better. So the placebo effect can be, in fact, quite powerful.

So when you now think overall about sources of utility, it's hard to imagine, in fact, any source of utility that's not influenced to some extent by beliefs. Often it's about how you look forward to some things. It's not just about going to a restaurant or a date or having fun with friends. A lot of it is also just the anticipation of all of that.

OK. So now more specifically, what do we mean by anticipatory utility? So many emotions in particular are intimately linked, related to what a person thinks about the future.

And you think about hope, fear, anxiety, savoring, et cetera-- a lot has to do with some stuff that might happen in the future, some events that might happen for sure, or some events that could happen but we're just worried about, such as when people are afraid and anxious about certain things. Those are actually things that might happen, even with a pretty low probability, yet people right now might be for sure quite anxious or worried or afraid.

Now then how do we define and think about anticipatory utility? Well it's utility derived now from anticipating in the future. So there's stuff that, in the future, enters your utility function directly. And those events that will happen in the future could be consumption, could be also services, could be like bad shocks or any things that will happen in the future.

Right now, in the current period, you're already feeling good or bad about it. Or something that will happen one or several periods in the future. And anticipatory utility is a prime example of utility from beliefs.

Now how might anticipatory utility affect behaviors? You can think about, broadly speaking, two classes of implications that we are going to think about. There are sort of other potential effects that might be there.

So the first one is people, if they have anticipatory utility, that might affect the choice of timing of certain outcomes. That is to say, if you have the choice when you can experience a certain good, you might want to delay that or not, depending on your anticipatory utility.

For example, if I asked you, would you like to go on a vacation next week or six months from now or six weeks from now? You might say, in fact, six months from now is preferable because then you'll have six months to look forward to that experience and be really excited about it.

Instead if I say, it's only next week, well then, it might be fun to go next week and maybe if you're discounting, you'd rather it earlier than later. But then there's not very much time to look forward to it and to savor and be excited about this experience. So anticipatory utility might affect the choice of timing of outcomes. We're going to talk about this in a bit.

Second, and perhaps more consequentially, anticipatory utility also may affect people's beliefs and their information acquisition. That is to say, if there is some things that I like to look forward to in life-- for example, if I think I'm a healthy person, I'm going to age really happily, and I'm going to live until age 90 happily, for a long time, well then for example, that anticipatory utility might then depend a lot on the beliefs about my health status.

And now if there were tests or some information that I could acquire about HIV, be it about cancer, be it about Huntington's Disease and the like-- it's very serious potential diseases, people might under-use or under-acquire such a type of information because they want to maintain overly positive beliefs about what's going to happen in the future.

And that could be potentially quite important because if people are not getting, for example, cancer screening, well then some treatments that perhaps could be done very early might get delayed, and that could be potentially quite costly. We're going to talk about this second.

OK. Let me show you some motivating evidence of anticipatory utility. This is not so much to establish the existence of such utility to you as like a very rigorous test and showing exactly this is anticipatory utility and how much of that is there, but rather to get you to start thinking about the patterns of behavior generated by it.

So George Loewenstein ingeniously asked undergrads about their hypothetical willingness to pay now to obtain or avoid certain experiences. So they asked them about the willingness to pay as a function of the amount of time until the experience occurs. So all values are normalized relative to students' willingness to pay to obtain the experience right away.

So you ask people about, like, how much are you going to pay for right away and 24 hours from now and three days from now, and a year from now, five years, 10 years from now? And then the willingness to pay right now is 1 by definition and everything is relative to that.

What are the experiences that you ask people about? There's sort of both pleasant and unpleasant experiences. One of them is receiving monetary gain or avoiding monetary losses. There's obtaining a kiss from a movie star of the student's choice. And there is avoiding a non-lethal but very painful electric shock. And now you can think about why Loewenstein asked about hypothetical willingness to pay, because implementing that would be pretty tough to actually do.

Now when you think about these types of experiences, you can think of it about like, what would your hypothetical willingness to pay be? And how would that look like as a function of the timing?

Would it go up or would it go down over time? And there's going to be two forces at play. There's going to be discounting at play. And there's going to be anticipatory utility at play.

We talked about discounting already a lot, which is, people like to have positive stuff in the present-- or rather in the present than in the future. People like to push negative stuff away to the future. So if bad stuff is going to happen, like an electric shock, you'd rather have that 10 years from now rather than right away.

So then the question now is, how does anticipatory utility affect those kinds of choices? And I already gave you some of the situation to start with, which is like, if you had a vacation that you potentially can experience, would you rather have this right now? Would you rather have it in a week from now, six months from now, or 10 years from now?

Well that depends on your anticipatory utility. You might prefer to have it in six months from now because that allows you to look forward to it for quite a long time. You might not want to have it 10 years from now. That might be good from the perspective of anticipatory utility, but it's really, really far away and you'd rather have things in the present than in the future.

So maybe if you take together anticipatory utility and discounting, you'd rather have it in an immediate amount of time from now, which gives you the sort of a form of an inverse U-shape of valuations over time, or valuations as a function of when-- right now of a function of what's going to happen in the future.

Let me show you what I mean. So when you look at Loewenstein's examples-- in particular, the kiss of the movie star is an example of an inverse U, where this is a pleasant experience. And what seems to be the case is that some people prefer that experience to happen in three hours, 24 hours, three days, and one year from now, compared to immediately.

You see the values here are always higher than 1, which is sort of what things are normalized to. So let me back up for a second. What does this graph show?

It shows the time delay on the x-axis, which is immediately, three hours, 24 hours, three days, one year, 10 years from now on the x-axis. That's when the actual experience is happening. And on the y-axis is the proportion of the current value. And it's normalized to 1, so the valuation immediately is 1.

Notice that these are willingness to pay to receive the experience or willingness to pay to avoid it. So everything is positive. Now for the kiss now, we see that the delay of three days from now seems to have the highest valuation. So people like to delay this for a bit. Notice that discounting would not generate you this pattern.

Discounting would say, you want it right away because you really enjoy this experience. So anticipatory utility would say, well three days from now and even a year from now it's better, because now you can savor it, really look forward to it, be excited about it when that's happening.

However, 10 years from now, people seem to like less than right now, presumably because it's really, really far away. So A-- I guess you're discounting that a lot and the experience itself gets discounted a lot. B-- even the anticipatory utility might only kick in nine years from now because it's really so far away that it's hard to think about this because it's so far away. And you might even experience anticipatory utility nine years from now. But you might discount that as well because it's really far away in the future.

So it seems like people have some preference for the present. There is some discounting going on. And there's sort of two forces at play-- discounting versus anticipatory utility.

Now in contrast, when you look at the negative shock, here it seems to be that people really seem to dislike having that shock in 10 years or one year from now, compared to having it right away. What's going on here?

Well if you have the shock in one year, 10 years from now, this is really an unpleasant thing that hangs over your head. So what, instead, people want is they'd rather sort of get it over with right away and not have to think about it too much overall.

Now and that's, again, sort of the opposite of discounting. If you just had discounting, what you would say is, well rather I'm willing to pay to avoid it if it's right away. It's really unpleasant right away. It's not so bad if it's 10 years from now.

So what discounting would say is, the function should be going down as you see it for losing \$4. And why is that? Well losing \$4 seems to be like, people rather do that in the future rather than in the present. And it's not really a thing that you think about it. You don't have a lot of anticipatory utility about losing \$4 in 10 years from now. That's not such a hugely important thing.

However, if you get an electric shock 10 years from now, that might be really unpleasant. So you really don't want to have that in the year or 10 years from now, having this hang over your head for a long time. So you'd rather have it right away to not have to suffer from any anticipatory utility, even though from a discounting perspective, that's actually sort of worse.

OK. And so I already said all of this-- subjects prefer a kiss from a movie star three days rather than now. They also prefer to have a shock now rather than in one year or 10 years. Remember this is willingness to pay to avoid a situation, so the willingness to pay right now for the shock right now is lower than the willingness to pay in one or 10 years. So the curve is going up.

Now this contradicts discounting of any kind. If you think about any discounting, positive or negative discounting-- as in, you might have a preference for the present or preference for the future-- you cannot generate these kinds of patterns. In particular, a pattern that's really hard to generate is a pattern where you get sort of these hump-shaped cases where, essentially, for inverse use, because either you want something in the present or really far away in the future.

But as I said before, if you had discounting, what you would get is essentially things are either increasing or decreasing. In the time horizon, positive things, you want right away. Negative things you'd rather have in the future. And this is not at all with Loewenstein finds.

Now what is the natural explanation? People look forward to the kiss. So they delay to enjoy anticipation. And it's very unpleasant to anticipate the shock, so they get it over with quickly. Now why did Loewenstein choose the kiss example? You can think about this for a second.

Well it is an experience with a high degree of savorability-- something you can think about and look forward to in the future. Loewenstein tries to rule out alternative explanations based on preparation. You might sort of say, well kissing a movie star right away might not be great because you want to get ready, shower, sleep well, and tell all your friends about it, and so on, or get some advice on which movie star to pick and so on. So there might be some reason that preparation might push you in three days from now.

Similarly, if he had said a dinner at a restaurant instead, you might say, well in a few days' delay is perfectly reasonable even without anticipatory utility. You might want to make a time right away. It might not be a good time. You want to get a date, et cetera, for the dinner. So choosing stuff later because of preparation seems totally reasonable.

That's not what Loewenstein is talking about here. He really is thinking about anticipatory utility. And so he has some other examples, but really, you can argue that really this is not about preparation, but part of that could be preparation as well. But what you really think about is going on is, really, this is about anticipatory utility.

When you think about, is preparation problematic for the electric shock in particular? Well the answer is no, because that would induce subjects to prefer it later, which is not what Loewenstein finds.

Remember, Loewenstein finds that people would like to have the shock right away, as opposed to a year or 10 years from now. And so preparation would say, well, I want it in three days and a year from now instead. But that's not really what Loewenstein finds. So maybe if you're not happy with the preparation explanation for the kiss of the movie star, preparation will not be able to explain to you the electric shock example.

OK. There's another example. So Loewenstein has this funny thing where he says, well getting electric shock has been a little bit of a weird example. So let me find something more realistic. And then being an academic, what he comes up with is cleaning hamster cages. So he asked subjects how much they'd have to be paid to clean 100 hamster cages.

I've never cleaned any hamster cages in my life. I imagine it'd be quite unpleasant. So you think about, what kinds of patterns do you expect to see? And what you expect to see is sort of something similar to the electric shock.

So people might really not look forward to cleaning hamster cages in a while from now and would rather do it right away to avoid anticipatory utility. That would be the pattern. So that willingness to pay for avoid it right away or you'd have to pay them not so much do it right away, compared to having to do it in a while from now, because the anticipatory utility makes things worse, as it did for the electric shock.

And this is what Loewenstein finds. If the payment is now for cleaning to be performed next week, it's going to be like \$30. But if the payment is now for cleaning to be performed in a year from now, the payment has to be higher. It's \$37. And again, discounting preparation or the like would create the opposite-- people would rather have it in the future than in the present, which is not what Loewenstein finds.

Notice, again, these are willingness to pay to have to do the experience. So how much would I have to pay you to do it?

Now again there's two forces here at play-- there is discounting versus anticipation. Now how do we think about discounting? So we talked about discounting before. And there are sort of three steps to thinking about discounting that we talked about. One was, there was a question like, what determines a person's instantaneous utility at each point in time?

We didn't talk about this very much during time preference. We were just saying like, here's an instantaneous utility function that determines how happy you are at each point in time. What we talked a lot about was, how do people integrate or aggregate those utilities across time?

So you have some utility today, tomorrow, two days, three days from now. And we then talked about, how do you aggregate, how much weight do you put on the different periods?

And then we talked about, what does the person predict about the future utility and behavior? Which is sort of the question about sophistication and naivete that we discussed as well. So as I said, the discounting issues you covered are about point number 2. The sophistication versus naivete issues are more about point number 3.

And here, people's prediction about the future utility and behavior was only relevant to the extent that it was helpful or harmful in helping people make good decisions. Right? So people didn't derive utility from being sophisticated or naive directly.

People only derive utility from their instantaneous utility function and from depending on the weight they put on those different instantaneous utilities. But their prediction affected their choices, and their choices, in turn, affected the utility. But it was never the case that people's predictions, their beliefs directly entered the utility function.

So instead, anticipatory utility is about point number 1. It's about the instantaneous utility function-- what goes into that, and perhaps people's beliefs about the future might do that. That's what we're talking about now.

Now notice that things like social preferences or reference-dependent preferences, those are also about number 1. Again, time preferences so far that we talked about were about items number 2 and 3.

Now what are the interactions between anticipation and discounting? We talked about this already. So you can distinguish between pleasant and unpleasant experiences.

So anticipatory utility is stronger for events that are closer in time, right? Something that's like 10 years from now, you're not that worried about right now. For something that's tomorrow or in two hours from now, you might be really worried about.

Now for pleasant and savorable experiences, some delay is optimal to have a few periods of anticipation. So think about each hour, each day as a period of anticipation. So you'd rather have three days than two days than one day of anticipation because each day counts as the day of anticipation, and you may value each day.

But you don't want it to be too far away because then there's also discounting. You'd rather have things at the present or close to the present than in the future. You don't want stuff to be in a year or two years from now because that's really far away, and then you discount that experience very much. Notice that you might also discount the anticipatory utility.

That is to say, suppose you have only three days of anticipatory utility of a nice restaurant or some nice experience. Only two or three days before that experience you're going to be excited about it. Well then if that's in 10 years from now, you're going to discount all of that. And then that's really far away. And if you really present bias or just discount a lot, then you're also going to discount the anticipatory utility.

So there's some trade-off for savorable, pleasant experiences between discounting and anticipatory utility, which might give you this inverse-U, hump shape power. Instead or in contrast, for unpleasant, fearful experiences, you want to either do it immediately to eliminate periods of anticipation-- so you might just say, get it out of the way. I don't want any anticipatory utility. And getting it out of the way essentially minimizes anticipatory utility because you get it over with.

Or you might put it off as much as possible for discounting reasons and to weaken anticipation. That's to say, let's just do it in 10 years from now. I might not worry about it for quite a while. I might also just discount any future anticipation.

And for discounting, if there's stuff in the future, I'd rather have bad stuff happen in the future than in the present because it essentially discounts the future. And it's good for me to discount bad stuff that's happening in the future.

So far we talked about the implications of anticipatory utility for the timing of consumption. Right? That is to say, when would you like a certain experience to happen? In the immediate present, in a few days, or further away in the future? But we took it as a given that something might happen and then the question was just when.

You could think of some examples where that might be important. Perhaps, more important overall for people's welfare decisions, and perhaps also for policy is people's information-gathering and beliefs. That is to say, you might get a situation where anticipatory or ego utility-- I'm going to show you-- might affect how much information people gather. And they might not be as informed as they could be otherwise because of those anticipatory utility reasons.

So one question you might ask yourself is, well, would an individual in a non-strategic setting with no anticipatory utility ever strictly prefer to refuse free information? And it's actually hard to come up with some situations where people just would refuse information if they're perfectly neoclassical.

You might say, well people might be overwhelmed. There's too much information available. And you might not be able to attend to everything.

But in general, I think the assumption in economics is, well the information could be useful for something. You might not even know that right now, but it might turn out that actually it's quite interesting information.

And in particular, usually, the assumption is there's free disposal. Right? If the information is not useful, just forget about it. Don't care about it. Whatever. Just write it down somewhere and whatever. It might be important in the future.

But really, why bother whether or not to have it? It's no problem. So you would never refuse it directly. You might not pay for it if you think it's useless, but you would never refuse to actually get it for free.

In contrast, somebody with anticipatory feeling cannot ignore information because that information affects his or her emotions. But if it's something that some piece of information makes you really unhappy, well it's very hard, actually, to forget about it and ignore it because once it's out there, once you've heard it, it's really that you can't unhear it and it's very hard to forget. In particular, it's hard to forget stuff that you really care about, that you're upset about or anxious about. And so it's really hard to do.

And so one quite interesting example is Dr. House and Thirteen. Thirteen is one of his employees. She's called Thirteen because I think he had 25 or something interns. And I think he fired almost all of them except for Thirteen who was number 13 of these interns.

He never-- at the beginning, at least-- bothered to learn people's names. So he just gave them numbers. So Thirteen was number 13. She was an excellent doctor.

It turns out that her family had a history of Huntington's disease. And Thirteen had the prime example, in the show, of information avoidance. She did not want to learn about whether she had the disease or not. And so the Huntington's Disease is such that, if your parents have the disease, your probability of having the disease yourself is way higher than in the general population. It's about 50% overall if one of your parents have the disease.

And so then you might think it's quite valuable to learn about whether you have disease or not, for a variety of decisions that you might make in the future or that are relevant for the future, ranging from when to retire, how much to exercise, what education to get, how much to save, whether you have children or not.

And Dr. House was usually pretty irrational. And there is choices or decisions that he does, was in fact being very neoclassical here and saying, you should get this information because it would be very helpful for you.

He, in fact, went behind her back, tested her, and sent her the results. And but then, Thirteen, for at least for quite a while, did not go to look at these results, presumably because she wanted to make herself think that she will be healthy and sort of delude herself in some ways about being a healthy person and not having this disease, even though she was not sure.

And so avoiding that information helps you do that because once you're tested, you're either positive or negative. And it's very hard to ignore that information once you have taken it. But before you have been tested, you could always sort of delude yourself and make yourself think that you don't have the disease.

Now what is Huntington's Disease? Let me be a little bit more specific. It's a degenerative neurological disorder. It's a very severe and heartbreaking disease.

Essentially, it affects the brain. It sets on around age 40 and really makes people very dysfunctional. And life expectancy is around 60 for somebody with Huntington's Disease compared to 80 or even higher for people who don't have Huntington's Disease.

People with parents or a parent with Huntington's Disease have about a 50% chance of developing Huntington's Disease themselves. Since the 1990s, there's a genetic blood test available. And so you can provide at-risk individuals with certainty about whether will develop it so you know for sure either you will develop it or you will not. And those lab tests cost about \$200 to \$300, plus consulting and other costs. The tests here are often covered by insurance, but most tests are paid out of pocket.

Emily Oster and co-authors of the paper that is on the reading list were arguing that people were not doing that partially to keep the test results private. Importantly, there's no cure or nothing that we can actually do to improve the disease. So when you think about other diseases, such as HIV or other like, usually there, there's a key motivation to learn about whether somebody is positive or negative, because then they can get antiretroviral treatment or the like.

Here that's not the case. So there's no actual cure or behaviors that you could do to that could mitigate the course of the disease. However, there's other potential behaviors that you could engage in-- so prepare you better and deal better with the disease once it sets on.

So what are these reasons? Why could such a test be valuable? Think about this for a second. There's a number of potential reasons that you might think about. I've listed a few.

There's a question about whether or not you want to have a child. There's issues about marriage. Would you like to get married or find a partner? Or do you want to tell your partner? And so on.

There's questions about retirement. When do you want to retire? For example, if you wanted to travel a lot after retirement. Well if the disease sets on at age 40, you might want to do that a lot earlier.

Education-- usually we think that the returns to education are high. So it's good news for anybody who is studying. But usually these returns are accrued over a long period of time. You're done with your education at the age 22 or 25.

You can then start working from age 25 to age 65 or the like or even longer. And so that's a long time. But if instead, you have Huntington's Disease, you might only have until age 40 or 45 or the like. So the returns to education are way lower.

You might also do things like participation in clinical research. Perhaps there's some chance that at some point there might be a cure. You maybe want to contribute to that. So there's lots of potential benefits of knowing, and that test could really be extremely valuable for you. And surely the value could be higher than \$200 or \$300.

Now what of the paper by Oster et al? So Oster et al. observe a sample of previously untested at-risk individuals over the course of 10 years. Why is it useful to have a sample of at-risk individuals? Well because the base rate, the rate of people who are not at risk, who don't have any parents with Huntington's Disease, that rate is quite low. So not being tested is not so much a puzzle.

If your chance is like 0.00-something percent, then it's really not that-- it's quite costly to do. It's \$200, \$300, plus hassle cost and so on. And there's many diseases one could have. Why test for Huntington's? It's perfectly reasonable to not test for it.

But here, there are people, these at-risk people, their chance is 50% to start with. So there's a lot to learn from. You go from 50% to either 0% of having it or 100% of having it. So there's huge changes in your beliefs, potentially.

While if you are a person who is not at risk, your chance of getting it to start with is like 0.00-something percent, with very high chance you're going to go to zero from that. So that's almost no change at all anyway. And there's only like a tiny, tiny chance that you're actually positive, in which, of course, you might want to change your behavior quite a bit. But the weight on that should be quite low because the probability of a positive test is really low.

So Oster et al find very low rates of this genetic testing. Fewer than 10% of individuals pursue predictive testing during the studies, over this long period of time. Many individuals get tested to confirm the disease rather than actually learning about whether they have it or not. You might think, what's really important is for the reasons that I mentioned here-- it's generating knowledge.

Right, you want to know, should you retire early or not? For that, you want knowledge. You want to reveal uncertainty. So if your chance to have the disease to start with is 50%, going from 50% to 0% or 50% to 100%, that's huge potential knowledge that's being created.

Once you think the chances like 99% or 99.9%, going from 99% to 100%, you actually don't learn that much. And you shouldn't change your behavior very much from confirming that you have the disease. Similarly, actually, there's similar evidence in studies on HIV, breast cancer, and other types of testing. People tend to systematically under-test even if the chance of having certain diseases is quite high.

And what's somewhat different for HIV and cancer-- any cancer, particularly breast cancer testing, there's at least some hope that if a disease gets detected early, one can potentially do something about it. There's some recent controversies or discussions on like, is cancer testing actually helpful?

And are we over-testing in the sense, is the test actually helpful in detecting potential diseases? But surely for HIV, it would be quite valuable if you knew early on whether you have a disease or not because then you could essentially take ARV, antiretroviral treatment, to help you.

Now let me show you first who's getting tested. So what you see on this graph is on the x-axis investigator evaluation of symptoms at the last visit. So on the left, these are people who are normal, who don't show any symptoms. On the right, people have like certain signs of Huntington's Disease.

And everything else in between, this is sort of in the middle of that. So further to the right means people have more symptoms. In particular, to the very right, these are people who have like certain symptoms of Huntington's Disease.

And what we see is, on the y-axis, whether people since their last visit get tested-- so there are several visits over time. And the raw means are what I wanted you to focus on. This is the upper of those two lines.

And what we see is essentially that testing rates are very, very low. The largest are about 5% in this particular sample. And that's essentially increasing in the symptoms that people have. So overall that fraction is high.

And in particular for people where you think the value of testing would be quite high, these are people who don't necessarily show any symptoms and people who have some symptoms but not very strong symptoms, you would think for them it would be particularly helpful to get tested because you could learn a lot. But instead, those are, in fact, the people who almost don't test at all.

The test rates are something like 1% to 3%. Now as I said before, for people with certain signs of Huntington's Disease, the objective knowledge from the test is very low since there's no cure. You're almost sure that you have it anyway, so in a way, you just really don't learn very much.

If you have these certain signs of Huntington's Disease plus you're genetically predisposed to have the disease, the chance of actually having disease is almost one anyway. So the test really does not give you a lot of information. However, perhaps, it changes your utilities from beliefs, which I'm going to talk about in a bit.

Notice that the perceived probabilities of having the disease might be lower even for people with certain symptoms. So what I'm showing you here is the investigator evaluation. This is essentially a doctor or neurologist actually looking at the symptoms and looking at, objectively, what's your chance of having the disease. That's different from what people actually think they have.

So I guess A-- and what we learned here-- is the testing rates are quite low. Second, people systematically under-predict or are overoptimistic about not having a disease. That is what I'm showing you here, is a motor score, which essentially is an assessment of how good are your motor skills.

And one unfortunate thing about Huntington's Disease is that your motor skills are deteriorating a lot. And so a high score means here, essentially-- if it's bad means essentially you're not doing well and you're more likely to have Huntington's Disease.

And what you see in the upper limit on the y-axis, is essentially the probability-- the objective and subjective probabilities-- of having the disease. The upper line here that you see at the top, that line is the actual probability, the estimated actual probability. And you see that's essentially increasing in the motor scores.

So people who have almost no symptoms overall, they have essentially a 50% chance of having the disease. That's the risk to start with. But then once people have more symptoms, that chance increases. And once you have a motor score in the 20s or 30s, your objective probability of having the disease is close to 1.

Now in contrast, the second line that you see here in the middle, that is the people's perceived probabilities. Those perceived probabilities are vastly lower and far less steep than the actual probabilities.

You see that people's perceived probabilities when they have no symptoms are still actually reasonably high, and they're pretty close to the truth. They're about 40%, compared to like something like 50%, 55%. So that's actually not that far off, perhaps because it's very clear.

If you have a parent who has the disease, everybody knows essentially, and particularly in this kind of study, everybody knows that the probability of having the disease is high, and so that perhaps, people can't really lie to themselves and know that the chances are pretty high if they have been untested. But then people seem to be largely in denial when it comes to these symptoms. People have these symptoms, and when you have scores of 20 to 30, their perceived probability of having the disease is really only slightly higher. It's about 50% or the like.

And so the people who are about here, well their actual probability is much higher. And people don't seem to update information that's really unfavorable toward them in terms of anticipatory utility because they seem to be feeling bad about it. Notice that there's actually a significant share of people who persist and really think about our reports saying that there's no chance of having the disease.

So there are some people who have a motor score of 20 to 30, and 10% to 20% of people here seem to say, I have a 0 chance of Huntington's Disease, which really seems to be unfortunately delusional in some ways. But really, these might be people who really just do not want to think about it and are just really worried and anxious. And they'd rather rule out that possibility of being sick for themselves.

Now one important question then is, well do individuals adjust their behaviors? Right? So far, we talked about, are people getting tested? And the answer is largely no.

Second, what are people's beliefs about the disease? And people seem to essentially be overly optimistic. Now there's a question, well, does that translate into their behaviors?

And now this is another little bit complicated graph. But essentially what the figure shows is coefficients relative to those tested negative for Huntington's Disease. Let's focus on the black graphs to start with.

These are people who are certain to have the disease, compared to people who are tested negative. So what we're comparing here is people who tested positive compared to people who tested negative.

And so the coefficients, if they're positive, essentially means that people are more likely to get divorced, more likely to retire, more likely to change their finances recently in their preparation. They're also more likely to be pregnant. So they change, essentially, their behavior relative to being negative.

So this is essentially just comparing positive and negative people. I should have said already, this is a type of setting where experiments are pretty difficult and ethically difficult to do, because if people really derive utility from information, providing people with information could really make them worse off.

It's quite different from other sort of experiments involving information [AUDIO OUT] wrong beliefs or sort of biased information about certain outcomes. Usually there, you would say, well let's provide them with good information. Maybe they improve their behavior. So there if you provide people with correct information, you could only make people better off.

Here, however, since people derive utility from this information, telling them the truth about stuff that they perhaps didn't know about could really make them worse off and deeply unhappy. And so one wants to be very careful with that. And therefore, there are not a lot of experiments in this space, at least for now.

So the study here that I'm showing you is essentially observational data. And in this case, comparing people who have tested positive, compared to people who tested negative. If you look at people who are just tested recently, arguably, these two types of people are not that different because when they made the choice of being tested, they didn't, of course, know whether they were positive or negative.

So that comparison really shows that people do adjust their behavior. So people who have been tested positive are more likely to change their lives in pretty dramatic ways. For example, people are more likely to get divorced, more likely retire, and so on.

People are also more likely to have a pregnancy. I think this might be for women only or for couples with their spouse. And that seems like an interesting result, in a sense. That's not what I would have expected.

Now that could be in part-- so the authors in the paper itself say that the sample that's used to look at pregnancies is relatively small. So maybe in some sense one needs to be a little cautious in interpreting those results. It could be that people try to get pregnant or have a child when they know about the disease because they want to have a child early on.

There seem to be some technologies where people can, in fact, avoid that their child had the disease. So you could make sure that doesn't happen. And there might be some couples who might say, well, even if it's the case that one of the parents will die early on, if the person was like 20 years old or the like, they might want to have the child early so they can spend sufficiently much time with the child before the onset of the disease.

But anyway, it seems to be that people-- the black lines are relatively clearly above zero. These are coefficients that compare essentially the black versus the group that's negative, which is the positive versus negative tested people.

Now when you look at the uncertain people, these are people who, essentially, with some probability, are positive and with some probability are negative. What we would sort of expect is that those people are somewhere in between the positive and the negative people.

So you would expect that the gray lines are also positive perhaps not as big as the black lines. Instead, you see essentially-- except for the pregnancy result, and that doesn't seem to be significant perhaps because the sample is small-- the gray lines are essentially pretty close to zero.

So essentially, people who are uncertain about the disease behave very similarly to people who are tested negative. So essentially, people who have not been tested or who are uncertain about the disease, they behave as if they were tested negatively and likely because they essentially delude themselves in some ways, thinking they are negative when there's no action that they need to take. So that really shows, or the authors argue that, essentially the lack of testing really changes people's behavior.

Notice that it's very hard for us to think about welfare-- are people better or worse off because of that? Because of course, it's a choice to do so. And you might say, well, I really value thinking about I'm pretty healthy and life is going to be good.

And that's really valuable to people. And they take into account that changes in their finance and so on would be valuable, but they're not doing this because if they were doing it, they would have to acknowledge to themselves that they are, potentially at least, sick.

OK. So how should we think about these beliefs and behaviors? Let me summarize what we have seen. So many people don't get tested, despite arguably good reasons to do so. People are often overoptimistic about the probability of not having the disease, despite often fairly clear signs that they have it.

And that really seems to be very much consistent with anticipatory utility. People want to feel good about themselves and about the future. And therefore, they're not getting tested and therefore they're overly optimistic.

Such over-optimism then translates into behavior. So people react less to the signs of likely having the disease than they arguably should. In particular, people who are not tested, they look a lot in their behavior similar to people who have been tested negative. And that's to say, perhaps if they were tested positively, they would change their behavior in ways that could be useful for them, at least along the road.

Now what models can explain such behaviors? It's important to understand neoclassical models would not generate these facts, because you would say, well, you should be tested because that would be useful. And utility is just not-- there's essentially no utility from beliefs. So why would you not get tested?

Let's say, well it's kind of expensive. But arguably, the value of doing so should be quite high. And the insurance would also pay for it if you were not able to afford it.

Now let me write down or show you a simple model. There will be a problem set, a question on this model for you to understand it a little better. So the model is extremely simple. It has two periods. And the relevant outcomes occur in Period 2.

In Period 2, the decision-maker will either be negative with probability p . So that positive person essentially would be healthy, with probably p of not having the disease. And that person will be positive with probability $1-p$.

The instantaneous utilities in Period 2, how they will feel. Think about Period 2 as being everything in the future. So Period 1 is the period when you think about potentially getting tested, which is right now or this month or the like. And Period 2 now is, say, next year, everything starting from all of the future.

Now the instantaneous utility of this long future period are u_{-} and u_{+} . u_{-} is if you're negative so u_{-} is the good case. And u_{+} is if you're positive. You sort of think that u_{-} is larger than u_{+} . It's better to be negative than positive in that disease or in general.

In Period 1, we're going to look at, now, people's choices. We can assume no discounting whatsoever. That's just for simplicity.

We also assume, for now at least, that nothing can be done about the person's condition. So A-- there's no cure. B-- there's also no financial or other adjustment that the person might be able to take to mitigate the onset of the disease. For now we're going to just rule that out.

This is just about, do you want to find out or not? And why or why not? Now what would I expect the utility theory to say? Here what is expected utility?

Well it's very simple. It's essentially, there's no anticipatory utility. Expected utility is just what's going to happen in the future. The probability p , you get u_{-} , which is the person is negative. With probability $1 - p$, the person is positive and then receives utility u_{+} .

Notice again, there's no discounting. This is in Period 2. But then since no discounting, δ or β are essentially just 1. Now let's add anticipatory utility to this model.

Notice that this type of utility can only depend on the person's beliefs about whether she will have the disease or not, not on what actually happens. Right? So the anticipatory utility depends on the probability of having the disease, and not on the actual consequences of the disease. And since the person in Period 1 has this anticipatory utility, doesn't know for sure whether she is going to have it or not, the anticipatory utility will only depend on the beliefs, on the probability p for perceived probability p of having the disease.

Now we're going to make two extreme assumptions about the formation of beliefs. Number one is beliefs are correct, in the sense that the decision-maker will always optimally update their beliefs. And once you've given the decision-maker a test, the person cannot just make up their beliefs in certain ways.

If the test is positive, the person knows they're positive. It's negative, a person knows that they're negative. So p is either 1 or 0 once they have the test. Or p is correct in the sense that the person knows it's p and there's no wrong information or beliefs about it.

Second, the decision-maker can choose to manipulate her beliefs. We're going to start with one, which is essentially the assumption that beliefs are correct. Let the anticipatory utility in Period 1 be f of p . And recall that p is the probability of being negative.

Now let's assume that, which is obvious in some ways and uncontroversial that u_{-} is larger than u_{+} . So essentially, again, remember u_{-} is the probability of being negative. So utility is higher in the second period if you're negative than when you're positive.

And then f of p , f is increasing in p . That's actually not so obvious in the sense it's not obvious that going from a certain percentage to a higher one, the anticipatory utility always increases. It's a reasonable assumption, but one might be arguing about it. But for now, we're going to just assume that f of p is increasing in p .

Now does the person want to know her status? Well let's think about this. So if she doesn't find out her status, her utility is f of p , which is her probability. Or without finding anything else is p . So f of p is just her anticipatory utility, plus the term that we had before, which is just the expected utility in the future in Period 2, which is p times u_{-} , plus $1 - p$, times u_{+} .

Now if she finds out her status, she might find out that she's either negative or positive. If she's negative, that's with probability p , her utility is the anticipatory part. If she's negative, she just gets f of 1. She knows for sure that she's negative. So there she is f of 1, plus u_{-} in the second period. All that, again, is with probability p .

If she's positive, her utility is f of 0. So her probability of being negative is 0. And then in the second period, you have u_{+} . Again, that is probability $1 - p$.

So now her expected utility, putting these things together, is the part here in the back, which is p times u_{-} , plus $1 - p$, times u_{+} . That's what we already had above. That's just the expected utility of Period 2.

And p times a negative test, which is f of 1, plus $1 - p$, a positive test, which is f of 0. OK. And so now, if you want to figure out, does she want to get tested? We just have to compare 1 versus $2 -$ or $2 - 1$ if you want, if that's the value of getting tested.

Notice that the expected utility part, what's going to happen in the future, will always be the same. We have $p u -$ minus, plus $1 - p$, $u +$. We have that here as well.

The reason for that is because we assumed previously that we cannot change the outcome. Right? So that's going to cancel out if we subtract $2 - 1$. So what we're going to be left with is essentially an expression that depends on f of p . p times f of 1, plus $1 - p$, f of 0.

When does she reject or seek information? Well she's information-averse if p times f of 1, plus $1 - p$, times f of 0 is smaller than f of p . That's exactly-- if you thought that's the case, it's exactly like the concavity condition.

If f is concave, that's in fact the definition of concavity. So if f is steeper for lower values, then the person will be information-averse. So if the person really dislikes any suspicion of bad news, for example, due to anxiety, there's not much value-added certainty, then the person will not get tested and not seek information.

In contrast, that function is convex, that is to say, if it's steeper for higher value, than the person really likes certainty. So some suspicion of bad news is not so painful.

That is to say, so essentially, when you think about p starting with something like 50%, what you want to compare is going from 50% to 100%, versus some 50% to 0. And going from 50% to 100% is more valuable. It improves your anticipatory utility more than going from 50% to 0% is sort of damaging of p of p -negative. It's a little confusing notation.

Then essentially the person will seek information. And it's the opposite if the person really dislikes any suspicion of bad news. And there's not much added value of certainty.

Now that is essentially just a question of, what does this f function look like? And we're going to try to figure out or try to estimate that function. Notice that it's a very simple example here, where essentially the person has to always-- there's no self-delusion or self-deception. The person always had to adhere to the truth.

Now consider instead the possibility that the person can manipulate their beliefs to make yourself feel better. Now one question you might ask yourself as well in the above framework, would you want to hold correct beliefs? Is there any good reason to hold correct beliefs?

And the answer is clearly no. Well there's no instrumental reason of beliefs. There's no cure. You can't really make anything better. You can't really adjust your behavior by assumption currently.

So why you would you ever think that you have to do this if you can make yourself feel better by increasing your p ? So if she could believe that she's HD-negative, for sure she would get higher utility for whatever happens later, which is f of 1 is larger than f of 0, or f of 1 is larger than any f of p .

Right? Essentially f of p is maximized. We assume that it's increasing. It's maximized that p equals 1. So really in the model that I showed you so far, if you could choose your beliefs, there's really no reason whatsoever of not choosing f of 1.

When there's a question, well why might you not want to choose f 1 anyway? And of course, the answer is, well that's because there might be instrumental reason. There might be actually some action that you might be able to take to adjust your beliefs.

So in particular, incorrect beliefs can lead to mistaken decisions. You can essentially make worse decisions if you have incorrect beliefs. So then there's a trade-off. Right?

There's a trade-off between, on the one hand, if you have incorrect beliefs, if you're overly optimistic, you feel good about it. And you're happier than you would be otherwise because you get to enjoy anticipatory utility. But of course, that might come at the cost of making wrong decisions.

So in that trade-off in general, decision-making with anticipatory utility, it's like overoptimism leads to higher utility than realism. That is to say that if you're fully optimizing to start with, some overoptimism would make you better off because essentially you get the value of overoptimism and then you're not going to distort your behavior very much if you were optimizing to start with.

So generally you can show that slight optimism, in many cases, in fact, an optimal thing to do, even if it comes at some cost. And of course, the intuition here is that the person wants to believe that she's healthy, and it makes her feel better. So she convinced herself that that's, in fact, the case. Now optimal expectations, essentially, as I said, trade-off-- anticipatory utility versus the value of knowledge for making correct choices.

And then, as I said, overly positive beliefs are an economic important implication of utility from beliefs. So once you assume or think about that there might be utility from beliefs, then one corollary is that people tend to have overly positive beliefs, because you explained that with utility from beliefs-- because really people have incentives to be overly optimistic even if that comes at the cost of potentially making worse choices.

Next I'm going to show you some other evidence of overly optimistic beliefs. And then in particular, we're going to talk about heuristics and biases, which is again the idea that people might be not good patients, in the sense that they might not be good at understanding probabilities and how to update probabilities accordingly.

Remember, recitation on Thursday and on Friday is going to talk about Bayesian learning. So if you want a refresher on that, that'll be quite helpful because on Monday we're going to talk about deviations from Bayesian learning. Thank you very much.