PRINCIPLES OF BEHAVIORAL ECONOMICS

Prof. Sujata Kar

Department of Management Studies
IIT Roorkee

Week 20

Lecture 20

Thank you. Hello everyone, this is the 20th lecture of the course on principles of behavioral economics, and in this module, we are going to discuss Bayesian updating and confirmation bias. So this is sort of a continuation of the previous module, where we introduced Bayes' rule along with several other basic probability rules or conditions, along with definitions. Now, what is Bayesian updating? Bayes' rule is often thought to capture how we should update our beliefs in light of new evidence. We update beliefs in light of new evidence all the time.

In everyday life, we update our belief that a particular presidential candidate will win the election in light of evidence about how well he or she is performing. The evidence here may include poll results, our judgments about his or her performance in presidential debates, and so on. This is also very much observed in the context of financial markets. We update our beliefs, our observations, based on the evidence on a daily basis. Accordingly, we make decisions about which stock to invest in.

Whether to invest in a particular stock or stay invested in a particular stock, whether we should withdraw our investment from a particular stock, and things like that. In science, we update our assessment about the plausibility of a hypothesis or theory in light of evidence, which may come from experiments, field studies, or other sources. Consider, for example, how a person's innocent belief that the Earth is flat might be updated in light of the fact that there are horizons, the fact that the Earth casts a circular shadow onto the moon during a lunar eclipse, and the fact that one can travel around the world. So with these further pieces of evidence, you can always update an individual's idea about whether the Earth is flat or round.

Philosophers of science talk about the confirmation of scientific theories. So the theory of how this is done is called confirmation theory. Bayes' rule plays a critical role in confirmation theory. To see how this works, think of the problem of belief updating as follows. What is at stake is whether a given hypothesis is true or false.

If the hypothesis is true, there is some probability that the evidence will be obtained. If the hypothesis is false, there is some other probability that the evidence will be obtained. So, in both cases, the probability of obtaining the evidence is actually going to change. Now, we will give an example.

Let H stand for the hypothesis and E for the evidence. The probability of H is denoted by Pr(H) and called the prior probability. So, this is the probability that we assign before the evidence comes into existence or we learn about the evidence. It is the probability that H is true before you learn whether E is true or not.

Probability of H given E is denoted by as we already know Pr(H given E). This is called the posterior probability. So, now you can see that once H comes into existence, you update your probability of H. That is why it is called the posterior probability. It is the probability that H obtains given that the evidence is true. This can be exemplified using a decision tree.

So, there are two possibilities: H is not true and H is true. So, we can further branch them out as Probability of evidence given H is true. Probability of no evidence given H holds. The probability of evidence given not H holds, and the probability of no evidence given no hypothesis.

Or not H holds, or the hypothesis is false here in these two cases. The hypothesis is true here in these cases. The question is what the posterior probability should be. This question is answered by a simple application of Bayes' rule, where we have already learned about this formula in the previous slide. This is called the Bayesian rule, where probability of H given E is expressed something like this.

$$\Pr(H|E) = \frac{\Pr(E|H)\times\Pr(H)}{\Pr(E|H)\times\Pr(H)+\Pr(E|\neg H)\times\Pr(\neg H)}$$

The result tells us how to update our belief in the hypothesis H in light of the evidence E. So, as and when E comes into picture, how we are updating our hypothesis. Specifically, Bayes's rule tells you that the probability you assign to H being true should go from Pr(H) to probability H to probability H given E. If you change your beliefs in accordance with Bayesian's rule, we say that you engage in Bayesian updating. So, moving from prior to posterior is basically the idea behind Bayesian updating. Suppose that John and Wes are arguing about whether a coin brought to class by a student has two heads or whether it is fair one imagine that there are no other possibilities that is the student is not willing to show the coin you are going to flip the coin and see the results but before that you cannot check whether the coin is a fair one or not you think that there is a possibility that the coin has heads on both sides and there is no tail on the coin For whatever reason, the student will not let them inspect the coin, but she will allow them to observe the outcome of coin flips. Let H be the hypothesis that the coin has two heads, so that not H means that the coin is fair.

It has one head and one tail. Let us consider John first. He thinks that the coin is unlikely to have two heads. His prior probability PRH is only 0.01. So, you understand that John's not H is 0.99.

Now, suppose the student flips the coin and it comes up heads. Let E mean the coin comes up heads. The problem is, What probability should John assign to H given that E is true? So, initially, this was his probability.

Alternatively, we can write this as his probability of H, having both sides heads. Now we are asking what this is when E is the observation that the first head has come up. Given Bayes' rule, computing John's posterior probability, which is the probability of H given E, is straightforward. We are given Pr(H) equals to 0.01. And therefore, we know that $P(\neg H)$ is 0.99.

From the description of the problem, we also know that the conditional probabilities, $Pr(E|H)$ is equal to 1 and $Pr(E|\neg H)$ is 0.05. This is 1 because we have seen that a head has come up. So, what is the evidence? The evidence is that if we go by the hypothesis that it has both sides heads, Then the evidence says that the probability is 1.

We have already got the evidence and it is given not H that one side is head and one side is tail. Then the probability of the evidence is 0.05 because in the next flip it may turn out to be a tail. All that remains is to plug the numbers into the theorem. So, we already know this formula. When we plug in the values, it gives a number like 0.02.

So, see, to begin with, John had probability H is equal to 0.01. Now, once the first evidence has come in, it has become 0.02, which is greater than 0.01. So, he is updating himself, his belief. Now the probability of the coin having both sides head is actually increasing. The fact that John's posterior probability H given E differs from his prior probability means that he updated his belief in light of the evidence.

The observation of heads increased his probability that the coin has two heads as it should. Notice how the posterior probability reflects both the prior probability and the evidence E. Now if John gets access to ever more evidence about the coin, that if he keeps on flipping the coin, then more and more evidence will come in. There is no reason why he should not update his belief again. Suppose that the student flips the coin a second time and gets heads again.

We can figure out what John's probability should be after observing this second flip by simply treating his old posterior probability as the new prior probability and applying Bayes's rule once more. So now it turns out to be 0.04. So again you can see that in the first trial it was, suppose I call it the first posterior probability and that was 0.02. this is the second posterior probability which turns out to be 0.04 greater than 0.02 so he is continuously updating himself about the possibility of the coin being an unfair one posterior probability increases even more after he learns that the coin came up heads the second time this is the observation from the second flip now

We consider a figure that illustrates John's and Wes's posterior probabilities develop as the evidence comes in. Notice that this is the diagram. So this is John and this is Wes. Notice that over time, both increase the probability assigned to the hypothesis. So, as the number of trials increases, their posterior probability increases.

This is the process of updating oneself. Also, that their respective probabilities get closer and closer. So, they are also merging and actually merging towards 1. So, after several trials, probability, posterior probability would be 1 if there is continuous You know, continuously heads keep on coming up.

As a result, over time, say after some 15 to 20 trials, they are in virtual agreement that the probability of the coin having two heads is almost 100%. So, which means the probability is one. So, this is the process of Bayesian updating. One striking feature of Bayesian updating is that John and Wes come to agree on the nature of the coin so quickly. After only about 15 flips of the coin, both assigned a probability of almost 100% to the possibility that the coin had two heads.

People sometimes refer to this phenomenon as the washing out of the priors. The priors are completely washed out. They have updated their posteriors to such a level that the prior does not have any impact anymore. That is, after so many flips, John and Wes will assign roughly the same probability to the hypothesis, independently of what their priors used to be. We have not talked about the priors of Wes here, but generally, the picture probably showed that he would start with somewhere like 0.5.

This represents a hopeful picture of human nature. When rational people are exposed to the same evidence over time, they come to agree, regardless of their starting point. So, even if two people start from very different points, evidence gradually convinces them to come to a common ground, to a common agreement. In real life, unfortunately, people do not generally come to agree over time, at least sometimes. Sometimes that is

because they are exposed to very different evidence. Conservatives tend to read conservative newspapers and blogs that present selected information because it supports conservative viewpoints. Liberals tend to read liberal newspapers and blogs that present selected information because it supports liberal viewpoints. So we generally read things and try to explore things that support our already existing viewpoints. Yet sometimes people have access to the very same evidence presented in the very same way as Wes and John do, but nevertheless fail to agree over time.

Why is this? So there are two possibilities. One possibility is that I continuously stick to certain things that support my viewpoint. As a result, I do not agree with individuals who have a different opinion. The second point is that even if there is evidence,

I still do not update my posterior or come to a convergence with other people with whom I initially disagreed. So, we fail to agree over time despite having several pieces of evidence. Why does this happen? So, this is something called confirmation bias. Part of the story is a phenomenon that psychologists call confirmation bias.

A tendency to interpret evidence as supporting prior beliefs to a greater extent than warranted. In one classic study, participants who favored or opposed the death penalty read an article containing ambiguous information about the advantages and disadvantages of the death penalty. Rather than coming to agree as a result of being exposed to the same information, both groups of people interpreted the information as supporting their beliefs. That is, after reading the article, those who were previously opposed to the death penalty were even more strongly opposed, and those who favored it were even more in favor. So basically, if some information is provided which is not very clear—either against any idea or in favor of any idea—

then even after getting that additional information, people may not change sides. They would stick to their own ideas and can become even stronger in their own beliefs. In the presence of confirmation bias, the following figure shows how people's beliefs change as they are exposed to evidence. Now, so far, we have considered three individuals. Individuals A, B, and C.

As usual, you are having a number of trials here—trials basically giving exposure to information to the individuals—and these are the posterior probabilities. Now, individuals A and B are behaving more like John and Wes, with trials are updating their posterior probability and converging to 1 but look at individual C he actually began from a particular point and stayed there was very close to 0, and his posterior probability does not change much—it actually rather converges to 0. So this is the case of confirmation bias, when the individual does not update his or her posterior probability despite having

a sufficient amount of information or evidence. In order for the evidence to be sufficiently convincing, we have individuals A and B, who updated their posteriors, but C did not update, which means he is suffering from confirmation bias. Confirmation bias can explain a whole range of phenomena. It can explain why racist and sexist stereotypes persist over time. A sexist may dismiss or downplay evidence suggesting that girls are good at math and men are able to care for children, but be very quick to pick up on any evidence that they are not.

A racist may not notice all the people of other races who work hard, feed their families, pay their taxes, and do good deeds. But pay a lot of attention to those who do not. So again and again, we see examples where we are trying to tell you that if people suffer from confirmation bias, they will pick up evidence that supports their existing beliefs. And that's why the posteriors are not updated.

Confirmation bias can also explain why people gamble. Many gamblers believe they can predict the outcome of the next game despite overwhelming evidence that they cannot. So why is this happening? Because if the gambler notices all the cases when he predicted the outcome correctly but fails to notice all the cases when he did not,

then of course he would focus on the situations where he predicted well, leading him to believe he will continue to do well. He forgets or ignores the evidence of his failures and continues gambling. So this is again a case of confirmation bias. The same line of thinking can explain why so many people think they can beat the stock market despite evidence that picking stocks randomly is just as effective. Finally, confirmation bias can explain how certain conspiracy theories survive despite overwhelming contradictory evidence.

The conspiracy theorist puts a lot of weight on morsels of evidence supporting the theory and dismisses all evidence undermining it. Scientists, by the way, are not immune to confirmation bias either. Philosopher of science Karl Popper noted how some scientists find data supporting their theories everywhere. So, he describes his encounter with Alfred Adler, the pioneering psychoanalyst. In Popper's words, "Once in 1919, I reported to him a case which to me did not seem particularly Adlerian,

but which he found no difficulty in analyzing in terms of his theory of inferiority feelings, although he had not even seen the child. Slightly shocked, I asked him how he could be so sure. "Because of my thousand-fold experience", he replied. Whereupon I could not help saying, "And with this new case, I suppose your experience has become a thousand and one-fold." So basically, this is a case of confirmation bias.

Just having seen the thousand cases—the thousand-fold experience he is referring to—he applies that thousand-fold experience to the thousand and one as well. Instead of checking that child or probably inspecting the child personally, he just went for a generalization, continuing with his belief related to similar cases. So that's why it becomes a case of confirmation bias. Popper's description makes it sound as though Adler is suffering from confirmation bias in a big way. Sadly, it is easy to find similar examples in economics as well.

Because of how easy it is to confirm just about any theory, Popper ended up arguing that the hallmark of a scientific theory was not the fact that it could be confirmed, but rather that it could, at least in principle, be falsified—shown to be false by empirical observation. Psychological research suggests that confirmation bias is due to a number of different factors. First, people sometimes fail to notice evidence that goes against their beliefs, whereas they quickly pick up on evidence that supports them. Second, when the evidence is vague or ambiguous and therefore admits of multiple interpretations, people tend to interpret it in such a way

their beliefs and reject all possibilities that would not. Third, people tend to apply a much higher standard of proof to evidence contradicting their beliefs than to evidence supporting them. With this, I come to the conclusion of this module, where we have talked about Bayesian updating, confirmation bias, and the application of Bayesian updating. Essentially, confirmation bias is the reason why Bayesian updating does not happen. In the next modules, we will carry forward with further discussions on applications of probability in the field of risk measurement, taking them up to the context of prospect theory.

Thank you.