

The following content is provided by MIT OpenCourseWare under a Creative Commons license. Additional information about our license, and MIT OpenCourseWare in general, is available at ocw.mit.edu. PROFESSOR:

Good afternoon. Before I forget, today's topic is cognition. And in part, a good chunk of this lecture is going to be based on work done by Kahneman and Tversky, who won the Nobel Prize for a bunch of this work some years back. A couple of years back, now.

The reason that is of particular interest is that Danny Kahneman is here tomorrow speaking. E-51, 4 o'clock, right? So if you happen to be a, interested in the material and b, at leisure at 4 o'clock tomorrow, I'd go hear him. He's a very good speaker and there aren't that many folks who I get to talk about in psychology who win Nobel Prizes. So, you should go hear him.

That's one characteristic of today's lecture. I realize that the other characteristic of today's lecture is that I get to tell you a bunch of things that are applicable to the forthcoming election. And to politics in general. This particular swath I'm cutting through, the large area of cognition, includes a bunch of topics that if you've watched the debates thus far, or if you go and watch the debate that's Friday, right, the next one's Friday evening, you'll be able to see some of the stuff I'm talking about today in action.

There are two large parts to what I want to talk about. And this is described, at least in rough detail, in this abstract at the top of the handout. We can make the distinction between the way, using the computer as a model for how the brain works is sort of a common game, rather like using the camera as a model for the eye. There are important differences.

One of the important differences that I'll talk about, that I started talking about last time, is that if you call up a document off your laptop, you are expecting that thing pulled out of memory to be exactly what you put into memory. As we saw, as we saw towards the end of the last lecture, that's not true about human memory. I'll elaborate on that a bit.

The other thing you expect is that your computer will do mathematical and logical operations correctly and quickly. And that also turns out to be not particularly characteristic of human cognitive behavior. For interesting reasons that I will describe in the second part of the lecture.

So let's start with this issue, this point, that says that recall is not just taking stuff out of storage. Last time, I talked about a couple of examples of that. And those are the ones labeled

as old on the handout. There's the context of remembering. If you are depressed, your memories of childhood are more gloomy than if you are in a good mood. There is retriever's bias, that was the example I was giving last time, with these experiments where the knife in a simulated altercation that you see, the knife might move from one person's hand to another person's hand in your memory, depending on the particular biases and associations in your memory.

Now, let me tell you about a few more ways in which memory is not just like pulling a document off the hard drive somewhere. One of these is what I've labeled on the handout as failures of source monitoring. Source monitoring is what we're busy making sure you do correctly on your papers. Telling us where the information comes from. And evaluating the information on the basis of where it comes from.

It turns out that that's not as easy to do as one might think it ought to be. Indeed, I got this quote from Euripedes on the handout that says man's most valuable trait is a judicious sense of what not to believe. If you hear something that's clearly false, is it possible for you to not be influenced by that.

This turns out to have been a question that philosophers have worried about. Descartes thought that the answer to the question was yes. If you hear a statement that's clearly false, you can label it as clearly false and, in a sense, put it aside. Spinoza, in contrast, thought that it would be very nice if that were true. But he doubted that it really was.

This has subsequently been tested experimentally. And I don't have enough colored chalk around, but I'll simulate it.

Here's what Dan Gilbert and his colleagues did a few years back. They did an experiment where you read little biographies of people. And the material you're reading is color-coded. And so you're reading along. And then some of it, here, we'll use straight lines for a different color. Some of it is red. If it's red it's a lie. Don't believe it. But if it's green, it's true. And if it's red it's a lie, and true, and so on. Read a paragraph of this stuff.

We're going to make sure that you've read it because we're going to test you afterwards and make sure that you actually know what it said. Because it's very boring to discover that you can ignore lies if you don't ever hear the lies. But we know you've heard these various lies.

And then what we'll do is ask you about this person. The metric they were using was a scale

between college student and criminal, I believe. And you were asked how criminal you thought these particular people.

So you're reading along. And so-and-so is taking 1803. And steals calculators. And does the problem sets reasonably reliably. And then fences the calculators. And so on, right?

The result is that even though you know that the statements are false, your opinion of this hypothetical person is lowered by negative lies. And, reciprocally, raised by positive lies. You can't completely discount the information that you're hearing, even if you know it's a lie.

Now, if it is not immediately obvious to you that this has direct relevance to the current political campaign, you can ask yourself about the behavior you can see in this campaign, or in any other campaign that you ever care to subject yourself to. Which is, you get charges that are made, typically by surrogates of the candidate. So a clear example, this here would be the charges about Kerry's Vietnam behavior, brought by the Swift Boat people, not directly by the Bush campaign. And I should be able to think of a flip side. Oh, the National Guard stuff for Bush. Regardless of the truth, we're not going to get into the truth assertion of this. But, in both cases the really nasty attacks are raised by surrogates.

If it is then proved to be a lie, then the candidate can have the great and noble opportunity to run around the country saying, look, people should not be saying that John Kerry molests small animals. It's a lie that he molests small animals. Anybody who says that should given a dope slap. And I would never say anything like, John Kerry molests small animals. And I will tell all of my people not to say that John Kerry molests small animals. And you can go home with this righteous feeling around you saying that I've done the right thing of condemning this smear. Not only that, I've spread it around a bit, and the dumb suckers are going to believe it. It's not going to swing the election one way the other, but it's going to push things. And political campaigns do this absolutely routinely. Perhaps not out of a completely Machiavellian understanding of the psychology, but because they know it works in some fashion. I don't know that they're sitting around reading Dan Gilbert's papers to figure this sort of thing out.

This has applications elsewhere in your life, too. It's very important to understand -- I'm going to take zinc because it's going to keep me from getting colds. Now, why do I think that? Is that because my aunt June, who's the medical wacko of the family thinks that's true? Or is that because I read it in the Journal of the American Medical Association? It makes a difference where the information comes from. So, some sort of an ability to evaluate the source is

important. And the fact that we're not very good at it is problematic.

A particular subset of the problems with source monitoring is what's known as [? crypt ?] amnesia. And that's the problem where you don't remember where your own ideas come from. The canonical form of this, in a setting like MIT -- actually, this is the part of the lecture for the TAs. The canonical version of this is you go to your doctoral adviser with a bright idea. Yeah, I got this cool idea for new experiment. Your doctoral adviser looks at you, says, that's stupid.

And you go away. You know, you're kind of devastated. Next week, in lab meeting. Your doctoral adviser has this really bright new idea. It's your idea. He thinks it's his idea. He does not, he's being a -- probably not, being a slimeball to say I'm going to steal their ideas. It's, you actually forget where it comes from. And in the absence of a clear memory of where the idea came from, if I've got an idea in my head about my research, gee, I wonder where that came from. It came from -- I think they come from my shower head. Because all the good ideas I have always come in the shower. But, most people tend to think it comes from them. I just cite the shower.

But this is actually a very good argument, by the way, for taking copious notes in a lab book when you work in a. Lab Because if it ever gets to the point of an argument, you want to be able to say, look, in this dated entry from October 5th, 10,000 and -- 2004, it might take you a long time to get your doctorate. here's the idea. I wrote down this idea, and then I wrote down this tear-stained comment where you told me I was a doofus. But, anyway, I really do think this was my idea. And I ought to be an author on the paper that you just submitted about it.

Anyway, it's a serious problem. Another place where that shows up is there are recurring scandals in the sort of academic book-publishing world where some well-respected author is proven to have taken a chunk out of somebody else's book. Or something very close to a chunk out of somebody else book, and publishing it.

Some people are doing that because they're sloppy, slimy people. Other people are doing it because you read a lot of stuff. Maybe you took good notes, maybe you didn't. And then when you start writing, this clever phrase comes to you -- or maybe just a pedestrian phrase comes to mind. But a phrase comes to mind. And you simply have forgotten the source of that phrase. Which is that you got it from somebody else's writing, and then -- it's a strong argument for making sure that you keep very clear track of where you're getting your material

from. And it's part of why we beat so hard on the notion of referencing in papers in a course like this. That's what [? crypt ?] amnesia is, in any case.

So, those are failures of source monitoring. It is also possible, to move on to the next one, to have memories that feel like memories but just are not true. Again, Liz Loftus has done some of the most interesting research on this.

One of her experiments, another one of these car crash experiments, remember the experiment of, how fast did the car smash, or bump, or tap, or whatever it was, into the other car? Well, here's another one. You watch another one of these videos of a car getting into an accident. And that part of the problem is that the car went straight through a yield sign. Later, you are given information that suggests that it might have been a stop sign. Can do this in a variety of ways, including ways that tell you that that is bad information. Another source monitoring problem. But, given later information that it was a stop sign, not a yield sign, people will incorrectly remember it -- not everybody, but you will get a percentage of people -- who will now misremember the original scene as having had a stop sign in it. If you ask what color was the sign, they'll say red, even though a yield sign would be yellow, and was yellow in the in the actual video that you saw.

So you've recreated, you've now put together the true memory with some false later information. And it's corrupted that true memory.

Now, that's of obviously some interest in terms of getting testimony in court and things like that. The place where this has become extremely controversial and a real problem, is the issue, is it possible to recover memories that were suppressed?

Suppose something really bad happens to you when you're a kid. Is it possible to suppress that memory and not have it available to you until some later point. At which point you now remember this memory. Is that possible?

The evidence suggests that it may indeed be possible to not think about something for a very long time. Have it bubble up into your memory.

Sadly, the evidence is also strong that it is possible for stuff to bubble up into your memory that isn't true. Under the right circumstances. If what's bubbling up into memory is memories of, for instance, childhood sexual abuse, for which there were a large number of criminal cases about a decade ago, now, it's a really important question to know whether these are

real memories. And it is deeply unclear that there's any nice, neat way to know the answer.

Now, look, obviously you can't, one of the reasons it's very unclear, is you can't do the real experiment. You can't take a little kid, abuse the little kid and say, let's come back 20 years and see if they recover the memory. I mean, it doesn't take a genius to figure out that that's not going to work. But, Liz Loftus has a very nice experiment that points in the direction of this sort of problem.

Here's what she did: you go and interview families about, let's take one specific example. Getting lost in a mall. Did so-and-so, the designated subject. You are --

AUDIENCE: Olga.

PROFESSOR: Olga. Did you ever get lost in the mall?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: OK, good. We won't ask her that. But we'll ask her family whether Olga ever got lost in the mall. And they'll say, no, no, no, it never happened. OK, good. Got a brother or sister?

AUDIENCE: Yeah.

PROFESSOR: Good. What have you got?

AUDIENCE: Brother.

PROFESSOR: Brother. So now we call her brother aside here, and say, we're going to do a little experiment here. Go ask Olga if she remembers the time that she got lost in the mall. And you do this with a lot of, Olga might look at her brother and say, no. And that's the end of that particular experiment. But some percentage of people, given brother who seems to think that this happened at some point, will, yeah, I remember that. Well, tell me about it. Well, you know how this story goes. Even if you were never lost in the mall, we could ask Olga. She could generate the story pretty successfully. It'll be something like, well, you know, I was there with my parents, right? And I was looking at these toys, like, in the window. And then I followed these, I was little, I was really little. So you don't get lost in the mall when you're 18 or something. So, I was little. I'm following these legs around. And it turns out to be the wrong legs. I thought it was my mom and it was somebody else. And I started to cry, and then this big man came and he took me to the security guys. And they announced it over the PA system,

that I was lost. And my mother came back -- people fill in really good, detailed stories of this. That are apparently invented out of -- well, not out of really whole cloth, I suppose. They're invented out of the understanding that you have of the possibility of being lost in the mall. And the childhood fear of being lost in the mall. But it's not a real memory.

Feels like a real memory. Subjects in Loftus's experiments, confronted later with the reality that this was all a setup, were often quite surprised -- Sure I was never lost in the mall? No, says Olga's brother. Liz Loftus just told me to tell you that. And, oh, ah, hm, that's a little weird.

But, the point is it is possible to generate things that feel like memories that are not memories.

Yes. AUDIENCE: What about

remembering dreams, supposing something happened in a dream that could happen, and you remember that as a [INAUDIBLE]

PROFESSOR: If you get routinely confused about what happens in your dreams and what happens in reality, that's known in the trade as a failure of reality testing. And is not considered a good psychiatric sign. But, sure, those sorts of things are possible. Particularly when you end up having desperately realistic kinds of dreams, the sorts of things where these, sort of, reality testing things do. Some dirt is, how many people, I'll probably ask this again when we talk about sleep and dreams. But, never mind. How many people have had the dream of their alarm clock going off? Dreamed that they turned off their alarm clock and got up. Only to then discover, oh, man, I'm in bed. And I'm supposed to be in my psych recitation or something like that. Or, at least tried to give that a story to their TA.

So, yeah, there are possible sources of confusion like that. The last item on this list of sources of possible confusion I think, says something like becoming plastic again. This is new stuff which I will come to momentarily.

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: That's OK.

AUDIENCE: Is there a way to distinguish between these false memories or if they've actually occurred? [UNINTELLIGIBLE] If you don't know them [UNINTELLIGIBLE] comes and says, I think I was sexually abused as a child, is there a way to determine, is this, like --

PROFESSOR: There is, at the present, time no ironclad way to know that. I mean, sometimes you can go and

get evidence -- there's no ironclad way to know simply from the memory testimony. There may be ways to find out by talking to other witnesses or something like that. There have been various efforts to look, for instance, at brain activity and see whether or not there's a difference in the pattern. So, you do this at the lab. Create a false memory and a real memory, and see if you can see the difference somehow in a fMRI scan or something of that sort. There are hints that maybe you can get there. It's like lie detectors. A fascinating topic for a whole lecture that I'm not giving.

There is no ironclad way of telling a lie from the truth by wiring anybody up. Regardless of what the FBI tells you. Actually, I shouldn't be telling you this, because the only reason that these things work is because people think they work. If you think that the lie detector's really going to work, then if you're lying, you tend to get kind of more nervous than if you don't think, than if you're telling the truth. And that's what they can monitor.

But a really good pathological liar, or a well-trained agent is apparently pretty good at getting past your average lie detector. Lots of people, you can get serious defense money by having a brilliant new idea about how to tell truth from lie and real memory from false memory. There's a lot of interest in doing that. But, no ironclad way of doing it at the present. Makes for great sci-fi novels, though, right what's going to happen when we finally figure out that we can actually stick this thing on your head and tell what your memories, whether your memories are true or not. But at the moment it's still sci-fi.

Not quite sci-fi, but still very new, is the indication. Remember, I talked about consolidation last time. This process of making a fragile memory into something that will last for years and years. There is evidence at least from rats in a fairly restricted range of studies, that when you recall a memory, when you bring it back out of, say, long-term memory into that working memory, that it becomes plastic again. And that perhaps some of the distortions in memory can arise at that point. Because the memory is again plastic. It needs to be reconsolidated in order to be, in effect, restored. So you bring it out. It's now vulnerable again in ways that it wasn't before you remembered it. My guess is that if you come back to the course ten years from now, that there'll be much more to be said about that. That's new and interesting.

But the general point should be clear enough. Which is that recall out of memory, retrieval out of memory, is an active reconstruction. It has multiple pitfalls in it that might cause what you are pulling out as a memory to be distorted in some fashion or other.

Now, the second part, the larger part, of what I want to talk about today is the sense in which we are not simple logic engines in the way that our computers, most of the time, are.

Now, one way to think about that is to think about a distinction between what's called narrative thought and propositional thought. Propositional thought, and this runs parallel. Narrative thought and episodic memory. Propositional thought and semantic memory. They travel together. So, propositional thought is the sort of thought that your computer would be good at. Moving symbols around, is x bigger than y , and problem set kind of thinking is mostly propositional thought. Narrative thought is thought that has a story-like aspect to it. Perhaps an episodic memory type aspect to it.

So, an example of narrative thought might be do you want to go to the party that is happening on Saturday night at Alpha Beta Gamma house, or something like that. How do you figure that out? Well you probably don't do some symbolic cost-benefit analysis. You probably imagine yourself there. Would I be having fun? Who's going to be there. You'd run a scenario. And that's the sort of narrative thought. Turns out that narrative thought has a power in our own decision-making that is capable of running across, running over, the simple facts of the situation. Even when we know the facts correctly.

Risk assessment is one example of this. So, forced choice. You gotta pick one of the other. Don't just sit on your hands here. Which makes you more nervous: taking a ride in the car or flying in an airplane? How many vote for riding in the car? How many vote for flying an airplane. As usual half the people can't figure out that forced choice, pick one or the other, means you have to raise your hand for one of these two options. At least if you're moderately alert. But anyway, plane wins big-time.

How many people, let's just ask the propositional thought, or the semantic memory nugget. Which is more dangerous? Flying in a car, that's only -- that's the Harry Potter version. Flying a plane or driving a car. How many vote for plane? How many vote for car? Everybody knows the proposition here. It's riskier to, you can measure this in all sorts of different ways. Per mile, per time. It's it's riskier traveling in a car than traveling in a plane. But people are much more nervous about planes.

Why is that? There are multiple reasons. But perhaps one of the driving ones is that the narrative surrounding planes, and particularly surrounding plane crashes is sufficiently vivid that it overrides your knowledge of the mere facts of the matter. If and when a plane sadly

goes down, that's big news. A lot of people tend to get killed. It makes big news in the papers and on TV. Highly salient. The fact that there's a steady drumbeat of car crashes, unless it happens to be very local to you, doesn't have any particular impact on you. And the raw statistic doesn't much matter.

In the same way, this turns out to be true in course catalog land. So that the MIT course evaluation guide, when it comes out, has vast piles of cool statistics. What people thought of this particular course. And then it tends -- I haven't seen the most recent one. But typically it has a paragraph of comments. One of my favorites over the years, so-and-so is to the teaching of physics as Einstein is to the National Football League. That's a very salient bit of narrative. And it turns out that those sort of comments have much more impact than all the statistics. In spite of the fact that the nice rational you knows that one person's brilliant catty comment, or brilliantly positive comment, is presumably much less valuable than the statistical average across 300 students. But the narrative, again, overwhelms.

You can watch this on Friday, when you watch the debate. You can be guaranteed that somebody is going to, it'll probably be either Kerry or Bush, actually. That Kerry or Bush will take some boring statistics and personalize it. They do this all the time. It's become a sort of it a trope, a standard item in the State of the Union message where the President always put a couple of heroes. They probably are, but, you know, heroes up in the balcony there so that he can point to them and say, look at this guy. Look at that guy. And in the debate it'll be, there'll be some discussion of employment statistics. And presumably this would be Kerry. Because he'll be beating on Bush about it. I don't care about what you say about the economy recovering or whatever. What I care about is poor little Olga. She got lost in the mall, and then her father lost her job as a mall security guy because of the decisions that your government, your administration made. And what am I going to say to poor little Olga. Why?

It's very sad about poor little Olga, who's now been lost in the mall and she's never going to sit in the front again. But there are, like, 260 million people, whatever it is in the country at the moment. It's very sad about poor little whoever it is. What's really important are the broad economic statistics. Broad economic statistics are wonderful propositional reasoning kinds of stuff. But are really boring. Unless you're an economist. But poor little Olga is really riveting. And poor little Olga is what you'll hear about in the debates.

We should do a debriefing about this on Monday. If anybody listens to the debate and hears any of these little tidbits, send me email. In case I miss it. And then we can we about this.

Now, this can get studied in the lab. Here, take a look. On the handout, there's a thing called the framing demo. Says that there's an unusual flu coming. It's expected to kill 600 people in the U.S., and there are two things you can do. There's Program A and Program B. Take a quick look, decide on Program A or Program B, and then we'll collect a little data and talk about it here. Don't cheat off your neighbor's test. OK. Everybody manage to come up with an answer here?

OK. So, let's see for starters. How many people went for A? How many people went for B? How many people don't know what they're going for? OK, well it looks about evenly divided. Boy, that's really boring.

Well, it would be really boring, except that there are two versions of the handout out there. So, Version One of the handout says that if Program A is adopted, uh-oh, I don't know which Version One says. Well, it doesn't matter which I call Version One, does it? Version One says, 400 people will die. Version Two says 200 people will be saved. Did I get the language right there? What you will notice it, if the net is 600 here, these are identical, right? There's no difference except in the way that it's phrased. Or the way the question is framed. There's a similar change to the second option, B in the die -- in this person, it's described as a $\frac{1}{3}$ chance that no-one will die and a $\frac{2}{3}$ chance that 600 hundred people will die. Do you know about the notion of expected value? You can calculate the expected value of that option, which will be $\frac{1}{3}$ chance that no-one will die and a $\frac{2}{3}$ chance that 600 people will die. And then that's going to be 400 people dying, again, here. So all the options are mathematically identical.

Now, how many people have Version One, the 400 hundred people will die thing? OK, how many people have the 200 people will die thing? OK, good. So, it's roughly evenly divided. Now, what I want to do is, I want to get the A and B for these two different versions. If you have the 400 people will die version, how many people opted for Option A? How many people opted for Option B? OK, so B beats A, and actually it beat A big time here. If you have the 200 people will be saved version, how many people went for A? How many people went for B? OK, in that case A beats B not quite as massively. But, clearly this manipulation has flipped the answer. It ain't math that's doing it, because the math is the same. What's changing here is what Kahneman and Tversky called the framing of the problem.

But really it's the narrative around the problem. What happens is, you read Option A. 400

people will die. That's terrible. Well, the second possibility in this case, the B is, woah, there's something we could do. That probabilistically might save some of those people. That sounds OK. I'll opt for this. As opposed to this one, 200 people will be saved. Well, that's better than nothing. And the alternative, B, here. You read the part that says, everybody might die. Hm, doesn't sound too good. And so you flip it.

So what you're doing here is the way that you ask the question changes the answer you get. Well, we can go back to the politics game here just fine. Pollsters do this all the time. The neutral pollsters desperately try not to do it. But campaign polls, when they're trying to get the correct answer out, who would you rather have for President, a strong leader who has ruled the free world with a firm hand for four years, or a flip-flopping senator from a liberal East Coast state?

Or you could ask the question, who would you like to have as your President for the next four years, a guy who barely made it out of Yale and has driven the country into a ditch, or a guy who actually got decent grades at Yale and is a national war hero and has really good hair? You can move the question around, by the way, and not just in polling. But the way you frame issues.

I mean, talk about framing issues. If you frame the war in Iraq as part of the War on Terror, and an integral part of the War on Terror asking whether you want to keep doing that, is different than taking the same thing and calling it a distraction from the War on Terror, for example. The facts remain sort of similar. It's harder to do the experiment in a nice, clean, controlled -- here, we can make all the math come out right. In a political debate, there's an argument about the actual facts, too, of course. But the way that you present the facts, the way you frame the facts, is designed to get the answer that you want out of those facts.

There's a beautiful round of this in today's paper, because there's the report of the whatever they're called, Iraq Commission that was looking for weapons and didn't find any weapons. And if you read the Kerry camp statements about what that means, and what the Bush camp says it means. The Bush camp is very big on intention this morning. He really, the report that Saddam wanted weapons. And so the Bush camp is busy saying, he wanted weapons and he would have gotten them as soon as he could. And the Kerry camp is looking at the same facts and saying, we all want stuff. Man, he didn't have anything.

So. Anyway, you can have lots of fun Friday. You can watch them do this one, too. All right, so.

The narrative that you put around a story. The narrative that you put around the story does not influence your computer. You can tell your computer about $\frac{1}{3}$ times 600 plus $\frac{2}{3}$ times minus 600, any way you like. You can use big font, you can use little font, you can use you want, it's going to come out with the same answer.

Not true of the way that you think about these things. You're going to get different answers depending on how you frame the question. Let me do a different example. That's not my handout, I got to see what's actually on the handout here.

So, right underneath that, underneath the bit about framing, there's another demo. I will, in advance, promise you that this is the same on all the handouts. Because otherwise you're going to sit there trying to figure out what's the weird thing. So, this is the same on all the handouts. I'm also aware that I have six slots between most probable and least probable. And only five options to go in there. Deal with it. Here's what you've got to do. You've got this guy Henry. Henry is a short, slim man. He likes to read poetry. He's been active in environmental and feminist causes. And Henry is a what? And so there are these five options here. What you want to do is rank order them from most likely to least likely. And then we'll talk about that a bit. Let's see how that -- No, you can't copy off your neighbor. You know Henry. I should have brought Henry in. Just grabbed some character and said, this is Henry.

All right. How many people need more time for this? OK, a couple of people still working on it. All right. Well, we'll just go up.

Let's look at a couple of particular comparisons on here. How many people put down, let's look at A and B. Whoops. That looks propositional thought. All gone. OK, let's look at A and B, for example. How many people said A was more probable than B? How many people said B was more probable than A? All right. B wins. And if I've got the right one here, so ivy beats truck big-time, right?

All right. How many truck drivers are there in the country?

AUDIENCE: A lot.

PROFESSOR: A lot. 10 to the -- I don't know the answer to this, but? 5? 4 is only a thousand.

AUDIENCE: 5.

PROFESSOR: 5 is probably-- it may be 6, but let's go with 5. OK. How many Ivy League classics professors

are there in the country? How many Ivy League schools are there?

AUDIENCE: Eight.

PROFESSOR: Eight to ten, something like that. Maybe, if you're lucky, there's ten classics professors at each. So that's going to get you to 10 to the 2nd. Is it a thousand times more likely that some random guy that you meet, who happens to be -- well, let's phrase this a different way. If you took all, whatever this is, a hundred thousand truck drivers in the country. I'm sure there are more than that. But let's just suppose there are only a hundred thousand truck drivers in the country. Are there fewer than a hundred of them who might be, for instance, short, poetry-loving feminists? And that assumes that every one of the Ivy League professors is a short -- there's no Ivy Leagues professors who are big, burly guys with tattoos and stuff. Or women, for that matter. Probably half of them are women, so this is already --

What you've done here is you've ignored what's known as the base rate. All else being equal, it's more likely that, if I just say, here's Henry, is it more likely that Henry is a truck driver or an Ivy League classics professor? It's way, way, way more likely that he's a truck driver. In fact, it's so much more likely that he's a truck driver that it's probably more likely. It's almost undoubtedly more likely. Even for this description of Henry the amazing feminist, whatever he was. You see, this is what's known as the base rate fallacy. People do not tend to take into account the base rate in the population when they're thinking about this. I'm asking how likely it is, and you're answering a different question. You're answering the question, how typical is Henry of your image, of your stereotype, of a truck driver versus of an Ivy League professor? And realize that even if -- let's suppose your stereotype is right. What's the chance that Henry could be a truck driver? Maybe it's only 1 in 100. All right, if it's 1 in 100, there's still 10 to the 3rd of them. And we're still swamping the population here.

So, failing to take into account the base rate is going to lead you to the wrong answer there's. Let's look at another piece of these data. Let's look at --

Is it more likely, A and E. A, E. How many people said that A was more likely than E? How many people said that E was more likely than A? Yeah, well this is this is Kahneman got the Nobel Prize, you see. Because he said that, too. Well, he didn't say that. He said, that's what you would say. Or he found, anyway. So E is much more likely. Remember, like, third grade set theory stuff? Let's make the little picture here. All right. The set of all truck drivers, right? The set of, I can't even remember what E is. E's like Mensa Ivy League truck drivers or

something. It's some small little set, it's a subset of A.

So, for starters it's a little unlikely that the chance of being in here could be greater than the chance of being in here. The best it's going to be is equal. In some fashion. Here's Henry. I'm meeting some random Henry. What's the chance that he's in E? It's just not -- Again, you're answering a question that says something about typical rather than probable. And you're ignoring what you would perfectly well know, of course, if we drew it out in set theory terms. You can't be more likely that something is in here, if the larger set includes that. That doesn't work really well.

So, people are very bad at base rates. Is this a bad -- is the purpose of this lecture to say we're all morons here? No, not really. The faults we have, the problems that we have in the way that we reason about problems like this, reflect -- What Kahneman and Tversky talked about were, if I'm spelling it right, heuristics. Sort of, mental shortcuts that do a certain amount of work for us. Our job isn't actually out there in the world to figure out set theory problems. Our problem is to get home safely at night. So, you're walking down the street at night. And it's a dark kind of street. And you see this big guy with a stick. And he's walking down the same side of the street as you. Do you cross the street?

Well, if you're in a Kahneman and Tversky kind of experiment you say, I know what I'm supposed to do here. I'm supposed to think hard about the base rates. Is it more likely, how many people are there in the world who are bad, evil, people with sticks that are going to, like, hit me. Versus, how many people are there who are, like, guys coming home from a softball game with their bat, or something like that? And they're nice people, mostly. Or at least even if they're bad, evil people, they're not really going to hit me.

The answer is, there's a very small population of people out there with sticks wanting to hit you. But, in the occasion that you get that wrong, it's a really bad mistake to make. And so, a mental shortcut that says not, what's the base rate here, but is this typical of a situation that could be dangerous, sort of, flipping it around. Could this be a bad thing. If I make a mistake and he's a really a nice person, so you might feel a little hurt that I crossed to the other side of the street. What's the big deal. If he's a guy with a bat and a nail in it and he's going to poke me in the head, and terrible things are going to happen. I'm not going to worry about the base rates here.

And your cognitive hardware seems to have been set up under those sort of constraints more

than it was set up under what might be optimal public policy solving constraints. So, this is why you can then run a political campaign where one side or the other spends a lot of time conjuring up images of bad guys with sticks, of some variety or other. Because you hear about the bad guy with the stick often enough. And if this guy promises he's going to save you from the bad guy with the stick, well, you'd better go with that. At least, that's an appealing thought. And it works for them. It works as a political technique.

All right, this is all -- who knows about Henry the, some weird truck driver that I invented or something like that? How about a realm where things really ought to be logical and mathematical. Which would be money. I mean, if there was ever going to be something that ought to just work out in terms of the math, you would think it would be money. And this is, of course, the reason why there is no Nobel Prize in psychology. But there is a Nobel Prize in economics. And that's what Kahneman, in fact, won. Kahneman and Tversky would have won it, but Tversky, unfortunately, died young. And you don't win the Nobel Prize posthumously.

So, Kahneman and Tversky also looked at the way that people reason about money. This is something that goes back long before Kahneman and Tversky. And, I can perhaps illustrate one of the older bits with an example. Let's suppose that you're going to buy a toy car for your cousin. You like your cousin, just in case you were wondering here. So you're buying this toy car for your cousin. And there are two toy cars out there at the moment. They are, for present purposes, functionally equivalent on all dimensions of kid loveliness or something. There's the midnight blue, I don't know, Model PT roadster or something from Toys R Us. And the midnight green Model PT roadster from Kids Are Toys, or something. So, the difference is the color. The green color is cooler. You know your cousin actually likes the green one better. It's a bit cooler.

Anyway, the midnight blue one costs \$12 and the midnight green one costs \$40. How many people are going to buy one of them. You have to pick one, again. How many pick the blue one? How many people pick the green one? OK. You guys have nice cousins, that's good. All right, the blue one wins big-time.

All right. You now have magically graduated from MIT. Congratulations. You got the job. You're now going to buy the car. And so, you're going to buy -- you're going to buy a real PT roadster. and And they come in two different colors. Amazingly, for today's purposes, they come in midnight blue and midnight green. The midnight blue one costs \$30,012. The midnight green one costs \$30,040. You like the green one better. How many people are going

to buy the green one? OK, so you graduated from MIT with this brain. It's the same. So, obviously it's now going very much in the other direction. It's the same \$28. What's your problem here?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: You were going to --

AUDIENCE: You're already paying over \$30,000 so the \$20 just seems very trivial.

PROFESSOR: Yeah. It's the same \$20. This is absolutely right, of course. But yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: The toy car is probably going to last you least as long as the roadster. I keep trying to fine-tune this to make this, you're not going to buy 15 cars here, and -- We're assuming a one-time purchase here, a one-time purchase here. And they're each going to last three years, or 100,000 miles, or whichever comes first, or something. Yeah.

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: That much more money? All right, all right. Would your answer have changed radically if I say, congratulations you have just graduated. It's also your cousin's birthday. You'd all of a sudden be -- Boy, there must be sort of a U-shaped function. Because by the time you get to a parent or something, ask your parents, who graduated a while ago. Hey, Mommy, can I have the midnight green car, it only costs three times what the other one cost.

Oh, actually I just sort of gave away the answer. It's the three times issue. And you were alluding to this already. That money is perceived, even though it's obviously the same \$28, that people think about money in ratio terms and not in absolute terms. This was actually picked up by Bernoulli a couple of hundred years ago. I don't know which Bernoulli. It turns out that there were buckets of Bernoullis. Five Bernoullis, and they were all geniuses. So, but one of the Bernoullis basically asked, what is the psychological value of money? Actually, in Kahneman and Tversky language, that's often called the utility of money as a function of actual money.

And what Bernoulli asserted was that the utility of money, the psychological value of the money, basically went up with log money. So it's basically preserving ratios and not the

absolute. It's not preserving the absolute values. Which, by the way, when you go out to buy that first car when you have all this money that you're apparently going to have when you graduate. The guy selling you the car knows all about this. And he knows that, do you want the plain vanilla one, or do you want the PT roadster with the really cool racing stripe on it. And this is part of our cool new graduate package, of stuff that cost us \$3.95 and we'll crank the price up by \$500 here, because on \$30,000, who can tell, right? And he knows. He knows that it's the same \$28 here and here. And his job is to get as many of those \$28 out of you as possible. And will take advantage of exactly this fact. If somebody said, I'll put a little stripe on your toy car for an extra \$28, well it's a smaller car. I'll put two stripes on it because it's a little car. You'll say, get away from me. But just listen to yourself go for the options when you go off and buy that car eventually.

All right. let us take a momentary break. And then I will say a word more about the oddities of money here.

[RANDOM NOISE]

OK. Let us contemplate a little more of the way in which the -- now, the reason this is important. The reason that this is the sort of stuff that eventually gets -- that where working this out is worth the attention of the Nobel committee in economics is that economists, for a long time, had this notion, had a psychological theory, in effect, of what people were doing economically. And it was the notion of a rational consumer, a rational person who's making decisions that were in a sense propositional. That they were doing the math to the best of their ability, and working things out on that basis. And this the Kahneman and Tversky program of research tells you is, the very interesting constraints on that rationality. The ways in which it sort of systematically deviates from what your rational computer might think about the same kind of question.

So, consider the following. Let's think about doing some gambling games. Where this is a one-shot game. You only get a chance to play this game once. I got a coin. It's a fair coin. If I flip it and it comes up heads, I give you a penny. If it comes up tails, you give me a penny. How many people are willing to play that game? Few people.

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: No, no. One shot. One time. The whole game is this one penny. How many people are willing to play one round of this game?

AUDIENCE: Give me the quarter instead.

PROFESSOR: Oh, she wants to play for real money. OK. That's fine. Let's do that instead. So. Most people are willing to play that, though. By this point in the lecture a variety of people are starting to wonder what is that I'm revealing about myself if I say I'll do this. So, the way you write that, you could write that as, you've got a 0.5 chance of winning a penny and a 0.5 chance of losing a penny. And the expected value of this is 0.5 -- you can figure that out, it's 0. It's a fair game.

OK, let's play it again. We'll flip, we'll play for her. Flip the coin. Heads, I give you \$100. Tails, you give me \$100. How many people want to play? Your computer doesn't care much about this, because the expected value remains 0, right? What's the problem?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: People are risk averse. Yeah, that's certainly one. That's what you find here. But why don't you once want to play?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: No, the expected value is just a thing defined in math land or statistics land. It's zero on this. But the two, you could either win or lose, it's true. You're not going to come out at 0. But if we played it for everybody in the classroom, if I had a lot of \$100 bills. Gee, I wouldn't want to play that game. Yeah. **AUDIENCE:** The expected utility of playing is less than the expected utility of not playing.

PROFESSOR: Hmm. That sounded cool. But I'm not being quick enough to figure this out. What did you --

AUDIENCE: I'm saying that winning \$100 is less good than losing \$100.

PROFESSOR: Thank you. That's exactly right. It turns out that this curve is not symmetrical around the origin. That negative money, the utility of negative money, tends to be somewhat steeper and more linear than positive money. With the result that, exactly what the gentleman said. That losing, that the absolute value of losing \$100 on this scale is greater than the absolute value of gaining \$100.

I mean, one way to intuit this is, there's plenty of room for you to gain \$100. You can do that all the time. Do you have \$100 handy that you can afford to lose? Maybe not. It would just hurt a

lot more. And if we decide -- And now we can -- There are some versions of this -- it's not that you'll never play a game for this kind of stakes. I mean, if I say I'll flip a coin. Heads, you give me a penny, tails, I give you \$100, how many people want to play? And the other ones didn't figure out what I said, presumably.

So, somewhere between that extreme and the even game is the point at which we would measure how risk averse you are. It'd be different for different people. I mean, would you play if heads, one time play, heads you give me \$50, tails I give you \$100? Now, some people would be willing to play. He's good, he's got \$50 handy. Or do you have a question? No, you were just --

AUDIENCE: [UNINTELLIGIBLE] 198.

PROFESSOR: 198, OK. He's got -- talk to him later. He's got more money than he knows what to do with. But, in any case, there will be a balance point at which your -- where the loss and the gain balance out. And then you'd be presumably willing to pay.

Strange things happen at the extremes. And that's the basis for a willingness to play lotteries. So, if you have a wager that says, heads you give me \$1, tails, I give you \$10 million, but the probability of coming up tails equal 10 to the minus, I don't know, 80 or something like that. I'm not getting the right numbers, particularly, for a state lottery. But it's not -- it's that kind of thing. The reason people run state lotteries is because they want your dollars. Not because they're interested in giving you a pile of money. So, even so by the time you get these very extreme kinds of gambles, as long as the cost is sort of minimal down here and the potential payoff is huge, for instance, you're willing to pay radically unfair games that you would never play in the range of, in the middle kind of range.

If I tell you, well all right, we can make this just a tenfold difference. So if we play a game that - - you can figure this out. At the extremes you're willing to risk \$1 to make \$100 million. Even if the odds are against you. In a more moderate game, you would not be willing to do that. And it's that sort of reasoning, or lack thereof, that allows lotteries to make their money. Even -- it's probably also confounded there by the fact it's not clear that the population as a whole really understands that this is a sucker's game. That the state is actually in the business of taking your money from you, not just redistributing the wealth. But even people who perfectly well understand that their chances are less than even of winning, even if they were to buy 10 to the 80th tickets, continue to buy lottery tickets and continue to. It's a strange way to raise money

for your state, but never mind.

So, money is like other aspects of reasoning. Is subject to the sorts of narrative stories that say, it hurts more to lose \$100 than it feels good to gain \$100. The sort of narrative that is beyond just the simple math of the situation. Let me -- OK, I have time to do one more example. Which is an interestingly illustrative one. And involves these cards. Which I believe I put on the handout. This is known as the Wason selection task, because good old Wason developed it. And, oh look, the rules are properly described on the handout.

You've got four cards here. Each card has a letter on one side and a number on the other side. You know that. What you want to know is, and here's the rule. The rule is, if the card has a vowel on one side, then it has an odd number on the other side. That's the rule that we want to check. Which cards, what's the minimum and full set of cards that you need to flip to check whether that rule pertains for this set of cards. That's the question. And so we'll vote on each card as we go along.

How many people want to flip E? All right. Lots and lots of people want to flip E. How many people want to flip the S? Hm. Couple of people want to flip the S. How many people want to flip the 7? A bunch of people want to flip the 7. How many people want to flip the 4? About the same bunch of people want to flip the 4. And some of the same -- I probably, well, it's too complicated to collect all the details of how many cards people think you need to flip.

The answer is, you need to flip two and only two cards here. And before I tell you which ones it is, I should say that like many of the demos in this lecture, this is the sort of thing that drives MIT students nuts. Because they -- it's a cognition lecture, it's a logic lecture. It's got, like, mathy things in it. And it's logic. And that's why I'm here. And I'm getting them all weird and wrong, and, and, and, and, and I'm going to flunk all my courses. And I'm going to die, or something. I mean.

I should note that one of the characteristics of the Wason selection task is that everybody gets it wrong. I don't remember if it's Wason in originally who did it. But he tried this out on logicians and they blow it. And if you're feeling, like, depressed after this, you take this to the calculus TA next hour, and see if she can do it.

So, as people correctly figure, you've got to pull this guy. Because if it's got an even number on the other side you're dead. And you don't have to do anything because we don't know -- there's no assertion about what non-vowels are.

You don't have to flip him. Because it doesn't say, it says -- here, let's do this in good logic terms. Which for mysterious reasons always talks about P. The assertion is, if P, then Q. Then odd. If P then Q. If you know Q, who cares. It doesn't say, if Q then P. So this could be an S without violating the rule. This guy, on the other hand, is not Q. This guy's not P. Everybody figured out we don't care about not-P. Not-Q is the other one that you do need to check. Because if this has an E on the other side, the statement is false. The rule is false. So you've got to check for an E on the other side of this.

So, there are now three people here who say, I've got it, I've got it. And good for you. That's nice. People routinely do very badly on this. And the interesting question is, is why -- well, actually, I'm not sure that's an interesting question. Why you do that. The interesting fact is that people do very well on other versions of exactly this problem. So let's set up another version of this problem.

We've got somebody with a -- a beer. We've got somebody with a soda. We've got an 18-year-old. And we've got a 22-year-old.

The rule is, you have to be 22 to drink. Well, 21. You have to be over 21 to drink. Which of these cards do we need to check? So do we need to check this one? Yeah, sure, right. If this sucker's an 18-year-old, he's going down. Do we need to check this one? We don't care about that. Do we need to check this one? Yeah. We need to check this one? No, we don't care about that. People are perfect at this, essentially. Sometimes if I phrase it right you can manage to get yourself a little confused. So if you got it wrong don't sit there and say, I'll never drink again.

So, people are perfectly good at this sort of real world example. It turns out that they're not good at all real world examples. I could, I won't bother trying to generate another real world example here. Leda Cosmides has argued that what people are good about reasoning about is cheating. Is somebody getting goodies that they're not supposed to Get Is this 18-year-old getting a drink when I can't get a drink yet? I'm going to keep good track of him. And the guy drinking this beer looks kind of young to me. Whereas, that four -- yeah, nobody cares about that. And if you reframe the problem, there are plenty of ways to reframe this, to put things like, oh, there's a version about, if you're in Seattle it must be raining. Examples like that. People bomb that all over the place. There's nothing at stake there.

Cosmides and a variety of the evolutionary psych people argue that your cognitive capabilities

have been shaped by the problems that you need to solve out there in the world. Is somebody getting goodies they shouldn't be getting? Is the guy with the stick going to hit me with the stick? That kind of thing.

It would be lovely if that was true. It might be true. The problem is that the research subsequent to this has made the issue more complicated. It is not -- there are cheating examples that people have cooked up that bomb. And there are non-cheating examples that people successfully manage to solve. But what does seem to be pretty clear is that we are not built to solve abstract problems. That's why you have to show up at MIT to learn that kind of stuff. Nobody goes to MIT to take courses in carding guys at bars. You come with that. You can figure that out for yourself. It's when you have -- unfortunately. The problem of running the country turns out to be probably a lot more like this than like this. But the decision is made like this. So, everybody go watch the debate and tell me what you find.