

## MITOCW | [watch?v=4RVTNoFrrYo](https://www.youtube.com/watch?v=4RVTNoFrrYo)

---

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at [ocw.mit.edu](https://ocw.mit.edu).

**ABBY NOYCE:** First, a little bit about how we make decision making. So, one model of how we make decisions is what's called the expected utility model. And the expected utility model basically says that you look at-- if you have two possibilities, you look at what you think the likely outcomes of each of these possibilities are, how probable those outcomes are, and how you value those outcomes.

So, if you think that if I make choice A, then I'm very likely to have a good thing happen but there's a small possibility of having something really bad happen, or choice B is kind of neutral, then depending on just how you value these different outcomes, you're going to make different choices.

**AUDIENCE:** Like expected value?

**ABBY NOYCE:** It's all about-- yeah-- expected value. So, you can think of this as like-- if you're trying to make decisions like what to do-- It's Friday night. You've got an exam on Monday. What are the things you could do?

So, you've got a decision. So, plan A is stay home and study, right? You could stay home. You could be a diligent student. And let's suppose that if you stay home and study, then it's almost certain that you'll do well on the exam.

**AUDIENCE:** There's a 1% chance you'll waste your whole time playing computer games.

**ABBY NOYCE:** Right, let's keep this a little bit simpler for a minute. Your other option is to do something else. What might you do on a Friday night if you don't want to stay home and study?

**AUDIENCE:** Go to a party.

**ABBY NOYCE:** Go to a party. Go to a party. What are some possible outcomes of going to a party? Good things that might happen, bad things that might happen.

**AUDIENCE:** You meet new people.

**AUDIENCE:** You might have fun.

**ABBY NOYCE:** You might have fun. Or it might be a party where you don't know anybody and nobody seems

to want to talk to you. It might be a party where you have no fun. So, what do you think-- give me a relative likelihood of these two things. Are these equally likely outcomes? Is one more likely than the other?

**AUDIENCE:** Well, it depends whose party.

**ABBY NOYCE:** Depends whose party. Think of a party. Pick one at random, some kid you know from school but you don't know really well. But they're having a party and you heard about it.

**AUDIENCE:** Were you invited?

**ABBY NOYCE:** Were you invited? I don't know.

**AUDIENCE:** We just go to the party anyway.

**ABBY NOYCE:** I've done that. So, all right. Let's say that there's a 75% chance that you'll have fun and a 25% chance that you'll be miserable. And we'll say that this is a certainty, that there's a 100% chance that you'll do well on your exam if you stay home and study.

How good or bad are these outcomes on a scale from like minus 10 to 10, where 0 is kind of neutral and minus 10 is the worst thing that could possibly happen to you and 10 is amazing? Where is having fun at a party on this value scale? Is it a 2? Is it an 8? Is it a minus 3?

**AUDIENCE:** 6

**ABBY NOYCE:** Probably not. A 6? Sure.

**AUDIENCE:** 7.5.

**ABBY NOYCE:** How bad is it to go to a party and be miserable?

**AUDIENCE:** 0.

**ABBY NOYCE:** It's just neutral? It's not bad?

**AUDIENCE:** That's a negative.

**AUDIENCE:** Minus 5.

**ABBY NOYCE:** Give me a negative. Someone give me a negative number. Minus--

**AUDIENCE:** Negative 7.

**ABBY NOYCE:** Negative 7? All right. How much of a good-- how important is it to you to do well on an exam?

**AUDIENCE:** 20.

**ABBY NOYCE:** Scale of minus 10 to 10.

**AUDIENCE:** 10.

**ABBY NOYCE:** It's a 10? This is super important? So, we probably want something that makes it more valuable than having fun at a party, maybe. Let's call it an 8.

OK. So, expected utility model would look at this and say, this option tree has a 100% chance of getting a value of 8. So, the value of this branch of my decision tree is 8. The expected value is 8. The value of this branch of my decision tree is going to be  $\frac{3}{4}$  but-- remember we assign probabilities to these. This outcome is a lot more likely. So, it would be 6 plus a quarter of minus 7, which is--

**AUDIENCE:** That's negative 7.

**ABBY NOYCE:** So, what is it? It's  $\frac{3}{4}$  of 6, which is 4 and  $\frac{1}{2}$ , plus a quarter of minus 7, which is minus-- so, this is 4 and  $\frac{1}{2}$ . This is minus 1.75, right? So 4 and  $\frac{1}{2}$ , 3 and  $\frac{1}{2}$ , 2 and--

**AUDIENCE:** .75

**ABBY NOYCE:** Thank you.  $\frac{3}{4}$ , yes. So, according to an expected utility model, if this is how you weight all of these things-- the goodness here, the badness over here and the probabilities of each of these-- then which of these is going to be the more rational outcome to pick?

**AUDIENCE:** Study.

**ABBY NOYCE:** Stay home and study. So, on the other hand, the fact that people probably-- that may not be the outcome that everybody would choose in this situation-- probably shows that people either assess these probabilities different or weight the outcomes differently. But this is a model of how you can expect people to behave.

**AUDIENCE:** I think that we forgot to take into account the being bored while we study.

**ABBY NOYCE:** So, you think that studying should be-- stay home and study--

**AUDIENCE:** It has a short-term outcome and a long-term outcome.

**ABBY NOYCE:** So, the total outcome should be maybe lower. You think our assigned values are off.

**AUDIENCE:** Yes.

**AUDIENCE:** What are we considering "well" on the exam?

**ABBY NOYCE:** Up to you.

**AUDIENCE:** 105%.

**ABBY NOYCE:** Sure. Anyway, this is a model. So, in cases where you think that people are going to make rational decisions, this is a model that gets used. So, like, economists think about decision-making in terms of maximizing your expected outcome and things like that. So, the thing is that decisions that are actually made by real live human beings, who are not perfectly rational creatures, as you've probably noticed at some point-- people make decisions that are different.

So, here's an example, kind of a hypothetical question that is a good example of how this works. Suppose the nation is preparing for the outbreak of a disease. We know that there is an epidemic coming. And it's expected to kill 200 people.

And the government and the national health people and the CDC are like, OK, there are two options that we could-- kind of like 24-- there are two options that we could do to deal with this. In one option, we are certain to save 200 of these 600 people. And in the other option, there's a one third probability that we'll save all of them, but there's a 2/3 probability that all 600 people will die from this disease. Which of these options would you support?

**AUDIENCE:** The first one.

[INTERPOSING VOICES]

**ABBY NOYCE:** You may not even be one of the 600 people who's expected to die in this. You may be safely ensconced in your, I don't know, your apocalyptic little hideaway. But who thinks that option one looks like a more appealing option here?

**AUDIENCE:** I'd run away.

**ABBY NOYCE:** Who thinks option two looks like a more appealing option? You are forced to choose. Pick one or the other. Option one? Option one.

Put your hands up where I can see them so I can count you. Make a decision. One, two, three, four, five, six. Option two? And three.

Now, look at our expected utility model. These come out the same if you count. So, OK, one third of 600, if you multiply the probability times the number, is 200 people.

So, according to this model, the expected utility of these is the same. And yet, a majority of people will like option one better. Why? Well, let's get to that.

Here's another scenario, same idea. We're preparing for the outbreak of a disease. There's been two options presented. The first option will result in 400 of these people dying, no matter what we do.

The second option-- well, there's a 33% chance that nobody will die at all, that everyone will be saved. And a 2/3 probability that 600 people are going to die. Which of these options do you like better?

**AUDIENCE:** Isn't that the same thing?

**ABBY NOYCE:** Who likes option one better? Raise your hands. Who likes option two better? Raise your hands. So, that's it.

All right, so by seeing them both together, you guys notice that these are exactly the same scenario, right? The numbers are identical. But if you present them to people separately, what we'll find is that in the first case, when this is phrased as how many people are going to be saved, people like certainty.

People are risk-averse. They don't want to pick a risky situation when the question is couched in terms of how much we're going to gain. When the question is couched in terms of how much you might lose, when we're looking at deaths, people start being more attracted to the riskier possibility.

When you have a chance of something bad happening, but a chance that it might not, people will go for risk. Whereas if we phrase it as a chance of something good happening or a chance that it might not, people will go for certainty. People like certainty for good things and risk for

bad things. We're risk-averse when things are discussed in terms of losses.

Does this make sense? Think of it as, like, gambling. If I tell you, you can either win \$50 for sure or you have a 75% chance of winning \$150 and a 25% chance of winning nothing, which of these are you going to pick? 50 certain dollars or a gamble on getting more than that?

**AUDIENCE:** More.

**ABBY NOYCE:** Who would pick 50 certain dollars? Who would pick the gamble?

**AUDIENCE:** A 75% chance of what?

**ABBY NOYCE:** \$150. Expected utility model says that these are equal, right? But when it's couched as things you can get, people tend to be-- and there's going to be variations in this for amounts and for personality-- people tend to be more attracted to the certain thing. If I spin that the other way around and I said, well you could either lose \$50 for certain or you could take a gamble where there's a 25% chance of losing nothing or a 75% chance of losing \$150, which one are you going to prefer?

**AUDIENCE:** The second one.

**ABBY NOYCE:** The gamble. Who thinks they prefer the gamble? Who thinks they prefer to lose \$50 for certain? Again, the expected utility math says that they're the same. But people like the risk.

**AUDIENCE:** Because there's a difference between losing \$50 and not gaining \$50. Losing \$50 means you lose money. Not gaining \$50 means you just don't-- nothing happens.

**ABBY NOYCE:** That's part of it. So, somebody put together something that's similar with a study, that they said, so, imagine you're at a casino and you either-- and they set it up as two possible scenarios. You win \$200 and you have a chance of gaining some more beyond that, or you win \$400 and you have a chance of losing some beyond that, so that all the scenarios kind of average out.

So, you ended up with \$250 or something. And I don't remember the exact numbers and I could probably work it out but it would take me 10 minutes. And what they found is that even when your net outcome would have been the same across either the win \$200 and then win some more versus win \$400 and then lose some, people are, again, more attracted-- they're-- a certain thing, when they have a chance of gaining something, they'd rather take a smaller,

certain gain, than a riskier, larger gain.

And the opposite is true when they stand to lose something. They'd rather take a riskier possible larger loss but possible not losing anything, versus a smaller certain loss. So, even when the amounts are-- so, our perception of the change is different. It's like we feel like gaining \$50 is less bad than losing \$50, that one of these is a bigger change than the other.

**AUDIENCE:** Is it like the [INAUDIBLE] when you have nothing to lose, like kind of go for it type thing?

**ABBY NOYCE:** It's probably related to that, although I'd say that that might more come from this risk attraction to risky-- when we're losing something. So, you're like, eh, what am I going to lose?

Actually I suspect that it's more like if you feel like you have nothing to lose, then you're not risking as much. So, what might seem like a risky possibility to someone who's got a vested interest just isn't as much of a risk. Like the bad things that could happen are less bad, in that case.

But it's probably related. There's a lot of studies looking at how people assess risk. And the outcome is that in a lot of ways we're not-- you can set up a lot of scenarios in which we're not very good at it.

Probably because the kinds of risks that we evolved to deal with, the risks you have to deal with as a primate running around on the Savannah trying to not get eaten by a lion, are very different from the kinds of risks you have to assess as a human being living in a city and dealing with cars and finances and all of these other things that we try and make sense out of every day. So, our intuitive math for these things doesn't deal well with a lot of situations that we find ourselves in. Yeah?

**AUDIENCE:** If you start off with \$200, don't people feel that, Oh, I already have \$200. So, I can only gain. I can't lose. Like, they have something could push off and go forward. Because if you have \$400, they're like, oh, I can lose more than \$200 or something.

**ABBY NOYCE:** Maybe, although I suspect you'd see the same outcomes no matter what your base amounts were. I suspect you'd see the same thing with \$10 and \$20 and values ending up somewhere in the middle. And I suspect you'd see it-- I think it's a factor of how the questions are phrased, whether it's set up as a situation of loss or a situation of gain, versus where the baselines actually are.

All right. So, humans are kind of generally bad at numbers. We're good at some other things too. So, the rules of thumb that people use when they're trying to decide-- and this is less decision making per se-- but when we're trying to figure out how to make sense out of the world.

Who here has taken some amount of CS, done some programming? One-- some amount of computer science, some amount of programming in school or for a hobby. A few, OK. So, when you're writing software, you're interested in algorithms, right? So, an algorithm is a series of steps that will always get you a right answer.

People, humans, human brains, don't seem to be so much on algorithms. We tend to take a lot of shortcuts. We tend to use what are called heuristics. And your heuristic is like a general strategy that will usually get you the right answer. It's good enough.

But if you're interested in how decision-making works, how people make some kinds of judgment calls or things like that, then-- just like how, in vision, it's interesting to look at the places where your visual system makes mistakes, looking at things like optical illusions in order to understand how it works-- people who study decision-making look at places where these heuristics seem to lead us to making bad decisions or decisions that are wrong relative to the facts or we will come to a conclusion that is not actually true.

So, I want to talk about three major heuristics that we seem to use in order to understand how our world works. First up is what's called the representative heuristic. So, suppose you toss a coin six times.

Which of these potential outcomes is most likely? Heads, heads, heads, tails, tails, tails, or tails, heads, tails, heads, tails, heads, or heads, tails, tails, heads, tails, heads, or heads, heads, heads, heads, heads, heads? Which of these is most likely?

**AUDIENCE:** It's the same, no matter what?

**ABBY NOYCE:** It is the same no matter what. How many people have kind of an intuitive perception that, for example, the last of these is less likely if you looked at that? Who would kind of, even if you know the math, know better from a math perspective. But it feels less likely, right? If you flipped a coin six times and got all heads you'd be like, whoa. And yet, it's no less likely than any other specific pattern.

But what happens is that because we know that we're taking a random sample from a pretty



random decision space, we think that it's more likely if the sample is similar to that population from which it's selected. So, this is kind of the, if it looks like a duck and it walks like a duck and it quacks like a duck, it's probably a duck, heuristic. If the sample that we're looking at seems to match the population that it's drawn from, then this is probably true.

So, representativeness heuristics, for all that we're going to talk about the places where they break, heuristics are generally useful. I mean, think about it. They've got to be. We wouldn't hang on to strategies for getting around in the world, if they weren't reasonably good at it. You'd get wiped out of the gene pool. It just wouldn't stick around.

So, representativeness heuristics are good for some things. They fail when dealing with randomness-- people are generally bad with randomness. They also fail when you're thinking about complicated categories-- categories that human beings tend to fall into.

So, if I was to describe for you my friend, Marie, and Marie is a sophomore in college. And she tends to wear a lot of black, like a black t-shirt with a long sleeve fishnet shirt underneath it, baggy black skater pants with some D-rings and straps on them. You guys know the look, right?

Which of these two things do you think is more likely to be true about Marie? Marie is an economics major or Marie is an economics major and one of the folks who puts on *Rocky Horror* at our university every Halloween. Option A or option B, which is more likely?

So, who thinks it's option A? Raise your hands. Who thinks it's option B? Raise your hands.

Which is odd because the math says that if I've got a conjunction of two things, where Marie is both an economics major and helps put on the *Rocky Horror Show*, that can't be more likely than simply the fact that she's an economics major. So, this is a case where the fact that Marie, this small sample of one, matches our perception of people who do *Rocky Horror* better than our perception of people who are economics majors. That makes it seem like one of these things is more likely. Something to watch out for.

So, this is, in part-- oh, one more one more kind of-- here, think about some stats questions. So, picture a town that has two hospitals, a big one and a little one. They both have an OBGYN unit. They both have babies.

About 45 babies a day are born in the big hospital and about 15 in the little one. And you guys

probably know that pretty much half of all babies are boys, right? But there's going to be some day to day fluctuation in this. So, you might have one day there's more boys than girls, one day there's more girls than boys. But over time, it would average out to 50%.

So, suppose that these hospitals are keeping a record of how often there were more than 60% of the babies born were boys. Which hospital do you think is going to have more days in a given year where there are more than 60% boys? More than 60% of the babies born that day were boys. Larger or smaller or about the same?

Who thinks it's larger? Smaller? About the same? That's what most people tend to say. But this is one of those things where you have what's called the small sample fallacy.

So, the sample size is-- the fact that the smaller hospital has a smaller sample size, it's going to see skews from the average more frequently. If I toss a coin three times, you probably wouldn't be shocked if they turned up all heads or all tails, right? If I tossed the coin 20 times, and it was all heads or all tails, you'd be a little bit more taken aback.

If I tossed it 100 times and got all heads, it would be really weird, right? Getting three heads for-- there's, what, like a 12.5% chance of getting three heads and something much smaller-- much, much much smaller than that-- chance of getting 20 heads, 1 over 2 to the 20.

So, when you take a small sample, the smaller your sample is relative to the population you're drawing from, then the more likely your sample is to be really different. Suppose I went to your high school and just closed my eyes, had the list of all the kids in your grade, closed my eyes, and picked three-- what are the chances they'd all be girls? Not huge but it wouldn't be shocking if I picked out three kids at random and they were all girls.

What if I picked out 50? It's a lot less likely. The larger the sample that we're taking, the more likely it is to be representative of the population we're drawing from.

So, one of the places where people make mistakes with this representative heuristic, is when they're looking at small sample sizes but not taking that into account when they try and consider what the larger population would be. This is the same mistake people make if they know three people who fall into a category and therefore make some kind of broad generalization, like, I know three guys who are engineers, therefore all engineers are guys. No, it doesn't work that way. But you've got to watch out for it.

All right. Another pattern people use is what's called the availability heuristic. So, read this list

of names. And look at me when you're done so I know.

Just read them. I'm not going to quiz you on them or anything, I promise. All right. Notice anything about them? Anything that calls itself to your attention? Yeah, some of them are authors. Good.

**AUDIENCE:** I only recognize the author names.

**ABBY NOYCE:** Here's a question. Were there more women's names or men's names?

**AUDIENCE:** Men's.

**AUDIENCE:** Men's.

**AUDIENCE:** I couldn't tell.

**ABBY NOYCE:** Who thinks there were more guys' names than women's names? Who thinks there were more women's names than men's names? Got to pick one. I'm going to make you pick one.

**AUDIENCE:** [INAUDIBLE]

**ABBY NOYCE:** Nope. Forced choice, you've got to pick one. Who thinks there were more women than men listed there?

**AUDIENCE:** [INAUDIBLE]

**ABBY NOYCE:** Who thinks there were more men than women listed? So, there's more guys than women. On the other hand, you guys probably noticed that the women's names are names you recognized, right? At least some of them. The guys' names, probably not. The guys' names are just kind of pulled at random.

So, the availability heuristic is a guideline we use, where we judge frequency of an occurrence based on how easy it is for us to pull up an example, pull up relevant examples of it. This is kind of the mistake that's being made-- who here knows somebody who is afraid of flying, doesn't like to fly, possibly because they think that they're going to die in some kind of terrible plane crash? Who here knows somebody who's afraid to fly, who drives or rides in cars that people drive?

What's with that? We all know that you're way more likely to die terribly in a car accident than a

plane accident, by like a factor of 100, right? But a lot of people feel like dying in a plane crash is more likely. And one of the reasons for this is that car accidents don't generally make the news.

Car accidents happen. They might show up on page four of the newspaper or something. Whereas plane crashes are a week-long media extravaganza. They're on the news 24/7 if a plane goes down. You've probably encountered this.

So, one of the things that happens is if you're trying to assess the relative frequency of car accidents versus plane crashes, then you might have an easier time retrieving the plane crash incidents from your memory. You heard more about them. You thought more about them. They were big news when they happened. And so, you perceive plane crashes as being more frequent than car crashes.

So, this availability heuristic is good for lots of things. If I ask you, do you think there are more guys or girls in your high school class, or is it about the same?

It's probably pretty close to about the same. But one of the ways that we answer this kind of question is by thinking through the members of our high school class that we know, and saying, how many of them are girls? How many of them are guys? And so, if the number of girls and guys that you can recall is roughly equal, you're going to guess that they're about the same. So, two things that affect how available something is in memory, are what's called-- one of them is recency, so, how recently we thought about something.

So, if you ask people, what do you think your odds are of dying in a plane crash, and ask people to estimate this-- and you ask them to do it and then-- well probably if there was a big plane crash and it's in the news, and a week later you ask people to estimate what their odds are of dying in a plane crash, they'll give you one number. And suppose six months later, you find the same people and ask them to give you another estimate, which of these numbers do you think is going to be bigger?

**AUDIENCE:** The first one?

**ABBY NOYCE:** Probably the first one, the one immediately after whatever it is. Because the recency of the event makes it more available to them. It's easier to remember. It's right there on top of your memory. And so, your perception of how frequently the stuff happens is skewed by that.

**AUDIENCE:** I heard that's how people choose the numbers that they gamble on when they're playing, say,

roulette. They say, hey, this number has appeared frequently. I'm going to bet on this one because it's a "hot" number.

**ABBY NOYCE:** Yeah, you'll see that. One of these things-- it's called the gambler's fallacy in something that's entirely random like roulette, like flipping a coin. Where-- but there's this kind of idea because it really feels like whatever has just happened is going to affect what happens next. Probably not true.

**AUDIENCE:** OK, is there a game where you have a bullet in the gun but it's only one bullet and then you spin the barrel and then you try to kill yourself?

**AUDIENCE:** I think that's Russian Roulette.

**ABBY NOYCE:** Russian Roulette.

**AUDIENCE:** It is a game, though.

**ABBY NOYCE:** I'm not sure I'd call it a game, but yes. More often you'll hear it in the context of being used as a metaphor, like such and so is equivalent to playing--

[SIDE CONVERSATION]

**ABBY NOYCE:** So, the other thing that affects this availability idea is how familiar you are with the outcome. So, again, and this is the thing where the level of-- how many people think that more kids get kidnapped in the United States today than 50 years ago?

**AUDIENCE:** I don't know. I've never really thought about it.

**ABBY NOYCE:** This is a more interesting question to ask grownups, people who actually were paying attention to the news 20, 30 years ago. Because the thing is that it's actually pretty stable-- kidnapping rates in this country are pretty stable. And yet, over time, you'll see things like parents get more concerned.

They don't let their kids play on their own as much. They don't let their kids get places on their own as much. You'll see these patterns, where if you ask your parents, were they allowed to go around town alone?

Were they allowed to take the T alone? Were they allowed to ride their bike to school? You'll see things where 30 years ago, people were allowed to do this stuff. And today, kids aren't.

And one of the things I think that affects this-- is and this is probably-- I'm going to try not to let this devolve into my "the news media is out to get you" rant, which is separate. But I think one of the things that affects this is that there's been this transition, so that things like, for example, a missing kid goes from being local news in the city where it happened in the immediate surroundings, to national news. So, our perception of how many of these are going on increases.

And because we know more about them, we are more familiar with the incidents. And so, this increased familiarity, again, makes it easier for us to think of them. We judge them as being more risky, more frequent occurrences. [INAUDIBLE]

One more heuristic that we tend to use, that I wanted to bring up-- this is what's called the anchoring and adjustment heuristic. So, when we're trying to estimate the size of something or figure out, I don't know, how much should it cost me to fly to Denver, we'll make an approximation. And then as we get new information, we'll adjust that approximation.

But what often tends to happen, is that that first number, that anchor number, we kind of treat it too much. Even if it was just a random guess, even if it had nothing to do, whatsoever, with whatever we're trying to estimate, it tends to kind of affect the changes that we make. We won't make big enough changes away from that original anchor.

So, one example of this would be, suppose I want to fly to Denver for this weekend-- which I do, but I'm not going to because I'm poor and Denver is expensive to get to because it takes a plane to it. Colorado. And suppose I said, OK, I think flying to Denver-- Suppose I set aside some money to fly to Denver, and I start looking for airplane tickets. And the first airplane tickets I look at are like \$1,200.

And I go, oh, boy, no not so much. And I start doing some hunting, some looking for cheaper fares. And, eventually, I hunt and I find smaller fares and smaller fares. And by the time I find fares that are \$700, it looks like a really good deal. So, I buy them because \$700 is a lot less than what I originally saw.

On the other hand, let's still say if I budgeted \$500 to do this, then \$700 is more than I had anticipated to spend. But because I had one idea set as this anchor of \$1,200, then adjustments from that feel significant, even when they're not.

Even if I hadn't still gotten down to what I'd planned to spend on tickets, it felt like it was a

much better deal than what I'd originally seen. So, this is the anchoring an adjustment heuristic. And so, what's interesting with this, is that people do this even when the information that seems to be set as an anchor has nothing, whatsoever, to do with what they're trying to estimate.

So, Tversky and Kahneman asked people to estimate percentages. What percentage of UN delegates are from Africa? What percentage of states in the United States have more than one state university? What percentage of-- probably a lot higher now than it was in 1974-- what percentage of, I don't know, people in the United States live in a city? Various percentages.

And they would ask the participant this question, and they had a wheel, like a mini Wheel of Fortune wheel, with just the numbers from 0 to 100 arranged in order around the rim. And so they would ask the question, they would spin the wheel, and then the wheel would stop on something. And then they'd have the participants indicate their answer by moving the wheel so that the pointer was on the percentage that they thought was true.

We got this? So, they've got a wheel. They spin the wheel. So, if I ask you, what percentage of UN delegates are from African nations. I spin the wheel, the wheel hits 10, you think it's probably higher than 10-- more than 10% of countries are in Africa. So, you're going to move it up from there.

And what's interesting-- what they found, is that this entirely randomly chosen number-- I spun this wheel right in front of you, you saw me do it you know it's random-- affects the outcomes that people give. You know, if wheel stops on a low number, most of the people are moving it up from that low number. The estimates that they give are lower than if it lands on a high number and people are moving it down, consistently.

**AUDIENCE:** Is there a set interval?

**ABBY NOYCE:** I don't know if they compared-- the way I would probably do this, is I would compare a group where I didn't have any wheel-spinning nonsense and just had them guess a percentage, and took the average of those, so that I could see, what does the general population believe the number of African nations in the UN to be, and compare from there. But I couldn't find any data that they had done this.

The example that the book in which I was reading about this experiment gave-- they gave an

example of the wheel landing on 10. Participants would give an answer of 25%. The wheel landing on 60, people had to move it down. They'd give an answer of more like 45%. So, a reasonably big range between these two but I don't have a good solid answer for it.

**AUDIENCE:** Could it be because they wouldn't want to move it so much from where it was randomly--

**ABBY NOYCE:** --because it was so heavy. Well, yeah. That's pretty much what we're saying, is that once a number has been set on, where-- not even consciously, but we're reluctant to move too far away from it. So, we tend to want to make small changes rather than big changes.

It's like, suppose your favorite ice cream place kept the cost of an ice cream cone at \$2 for years and years and years. And then all of a sudden, it went up to \$4. You'd be like, what? The price of my ice cream just doubled, I don't think so.

**AUDIENCE:** Ridiculous!

**ABBY NOYCE:** You'd be ticked off, right? You'd probably buy less ice cream. But if it went up gradually over the same amount of time, like, instead of being the same for so long and then going up suddenly, if it went up by \$0.25 a year or something, you'd be a lot less averse that.

So, that's another example where big changes in some kind of numbers are disconcerting. We don't like them. Little changes we're pretty OK with.

All right. So, kind of take-away things from this. So, decisions that we make, judgments that we make, all of this are influenced by a lot of things that we aren't directly aware of. One of the things that's interesting in all of this heuristic study, is when you point out to people that people tend to use these heuristics, that they tend to make these kinds of statistical decisions, that this is a piece of information you should take into account, they get better at it.

So, there's at least-- so, we use these heuristics and kind of naively we don't think about what the flaws in them are, but people can learn to take that into account and make better decisions and more accurate estimates by being aware of how their heuristic systems work. OK.

Shifting gears a little bit, I want to talk about one particular study. So, we were talking a minute ago-- for the first part of this-- about the strategies that people use to make decisions in their life. These guys are looking at what some of the underlying neuroscience is in this. And they had people do a very simple task.



They said, we're going to have you have a button in each hand and whenever you feel like it, push one of them. So, they had people do this. And they had them sit-in an FMRI tube. And they wanted to know, is there any way that we can measure what kind of decision people are going to take before they're aware that they've decided it?

So, here you are. You're sitting in the FMRI tube and you're seeing a screen which is showing a stream of letters for 500 milliseconds, for half a second, each. So, it'd show a letter, show another letter, show another letter, show another letter. And you're just hanging out there. And whenever you feel like it, you are asked to do two things-- to press either the left or the right button and also to make a note of which letter was being shown at the moment that you made the decision.

So, you were asked to decide and press right away-- not to decide and then wait for a minute, just make a decision, push the button, and notice which letter was being shown. So, what they did is after you pushed the button, you got a screen and you were asked to indicate which letter had been displayed when you had made the decision to press the button. So, if I'm sitting here watching the screens and in this case, when the moment when the Q is shown, is when I decide to press the button. And no matter how fast I go from decision to button press-- it might be a little bit slow-- so what's actually being displayed when the computer detects the button press, might be the V.

And after I press the button, the computer's going to ask me, of these three previous slides, which one was showing when you decided to press the button? And if the Q was showing when I decided to press the button, I'd indicate that. I'd say it was the Q and move on and do it again.

So, is what they're asking subjects to do make sense? Decide to push a button, remember which letter it was, tell us which letter it was. And what these guys wanted to find out was a couple of things.

There's been research in the literature for a while, for about 20 years now, showing that when people decide to make a motor movement, like pressing a button-- which is a motor response; it's a very small one compared to a lot of other motor things-- then at least some period in advance of that, maybe half a second in advance of that, there's activity in the brain that seems to indicate that a decision has been made before the participant is aware that they've made a decision. Before you say, I've decided, before you can think that, your brain

apparently already knows that you've decided. Creepy, yeah.

So, these guys said, OK, the previous rounds of this had only asked people to push a single button. They didn't have a choice. So, we want to see if we can predict the choice that people make.

Here they've got two options. You can push either the button in your left hand or the button in your right hand. And they want to see if they can predict which choice people make before the people are aware that they've done it.

So, what did they find? Look, pretty brain pictures. OK. So, they had people do this in an fMRI. They looked at their patterns of brain activity in particular regions over a temporal scale.

So, each of these-- as you can see, there's different particular brain regions that are being graphed here. Each of these graphs has time along the bottom axis in seconds. This red line is the point at which subjects reported that they had decided to make a decision using that-- which letter was showing when you decided-- that method. And then this is showing how accurately the vertical axis is showing.

So, they looked at the particular patterns of activation in each of these regions. And they said, can we find a pattern here that predicts-- Using the information in this brain region, at this point in time, can we predict which choice you made? So, this is the 50% line here. So, that's chance. And then anything up above that-- so, this would be like 75% accuracy. That's pretty good.

So, the filled-in black dots are places where they're significantly better than chance at determining which choice was made. The white ones are places where they're not, where it's just the same as chance. So, what did they find out?

Well, they found out that after subjects made this conscious decision to push the button, then all of these motor cortex areas are really good at predicting whether you're going to press it with the left hand or the right hand, which isn't too surprising. Motor cortex is involved in all of these kinds of decisions.

Left motor cortex, right motor cortex, the supplemental motor area and the pre-supplemental motor area, which are, again, as you guys might remember from last week, are areas that seem to deal with planning activities and goal-directed motion and stuff. And so, all of these are areas that can identify what choice you're making after you've consciously made it. What's

more interesting is these ones, which show some amount of ability to predict what choice you make before you're consciously aware that you made it.

So, even way out here-- look at this. This is frontopolar cortex, lateral and medial. So, way out here. Eight seconds before participants report that they've made a decision, just looking at activity in the frontopolar cortex, these researchers could predict with 60% accuracy whether you were going to push the right button or the left button. Eight seconds-- that's like forever in neuroscience terms-- before you know that you made that decision. I can tell you which decision you're going to make.

**AUDIENCE:** That's creepy.

**ABBY NOYCE:** Yeah. Slightly sooner in posterior cingulate cortex. It's not quite as good in medial frontopolar. But they said if they combined the information from all three of these areas then their accuracy actually went up. It went up above that 60% to like 75%, 3/4 accuracy at predicting whether subjects were going to push the left or the right button before they thought they'd made that decision, which is cool.

**AUDIENCE:** Well, couldn't they have a 50% chance of [INAUDIBLE]?

**ABBY NOYCE:** Right, so it's getting things above 50% that are relevant. So, again, the filled- incircles on the graphs are time points at which the difference between their accuracy and chance is statistically significant, which-- it's not just random that they got this, that it actually means something. Whereas the white ones, where it's very close to chance, are not significant. Questions, comments, creepy?

**AUDIENCE:** I think there was an article in, like, *The Wall Street Journal* about that.

**ABBY NOYCE:** Yeah, there's been results like this for quite a while. It's kind of an ongoing topic of research. Like I said, the first people doing this were doing it in the '80s, in the mid '80s.

But as the methods get better and as new techniques get looked at-- so I think this is the first one like this where they have the two-choice option. This is a paper from like May of this year, I want to say. This is new. This is shiny new.

**AUDIENCE:** What happens if you visualize pressing with your left hand as you do it with your right?

**ABBY NOYCE:** I have no idea. You'd probably mess up their data. Only do this to researchers with whom you

are not friends, at least without telling them that's what you're doing.

**AUDIENCE:** That's a good test to do [INAUDIBLE]

**ABBY NOYCE:** Maybe. Yeah. And on the other hand, we know that when you visualize doing something, at least to some extent, the pieces of cortex that would actually be involved in doing it, are active. So, if I ask you to visualize a banana, the same parts of your cortex that would be active if you were actually looking at a, banana are active. Or if I ask you to imagine standing up and doing a jumping jack, then the pieces of motor cortex that would be involved in that, are active.

All right. So, this brings us around to some of the big philosophical questions in this field. So, we've been talking pretty much all through this course with this underlying assumption that our subjective experience of the world-- our perceptions, our understandings, our memories, all this-- can be traced directly to physical stuff going on in the brain. And we're pretty sure that physical stuff has to be caused by other physical stuff.

You can't have something that is only mental without a physiological basis that can then cause physical stuff. So, physical changes in the brain have to be caused by other physical changes. Some of the stuff where we've looked at the chemistry or the cellular-level changes underlying things-- that's this idea that physical changes are caused by other physical events.

And this starts bringing us into all of these fun questions about free will. So, if you decide to do something based on the state your brain is in right now and the state your brain is in right now can be traced back to every state your brain has always been in then do you get to make a free decision or are all of your decisions predetermined by the experiences you've previously had?

**AUDIENCE:** This class has made me stop trusting everything.

**AUDIENCE:** That's a logical argument.

**ABBY NOYCE:** I'm OK with that as an outcome.

**AUDIENCE:** I would say no. But I would say both answers could be arguments because if you said yes, you could think that it's-- I don't know.

**ABBY NOYCE:** So here's a question. Suppose you are on the jury in a criminal case for a guy who's been accused of beating up a bunch of people, like a serial assault, like assault on a number of

people, has injured people. And what his attorney says is that this guy-- you get lots of witnesses saying that up until five years ago this guy was gentle as a lamb, never hurt anybody in his entire life. And his attorney says that he got a brain tumor and he now has a brain tumor that has caused this change in his behavior. Is this guy responsible for what he has done? Should he be convicted?

**AUDIENCE:** Yes.

**AUDIENCE:** No.

**AUDIENCE:** Yes.

**AUDIENCE:** No.

**AUDIENCE:** Yes.

**AUDIENCE:** --because he got a brain tumor, then it's not his fault. He didn't choose to get a brain tumor.

**AUDIENCE:** Yeah, but he still hurt other people.

**ABBY NOYCE:** OK. What happens if we say, OK, you have to be treated for this. We remove the brain tumor, and this guy goes back to his previous exemplary, gentle-as-a-lamb lifestyle? We postponed the trial for a year for him to get this treatment.

Then what? You going to throw this guy in jail? He's got kids. He's got family. He's got a good job. He's a staple of his community. But he beat up 10 people and put them in the hospital.

These aren't easy questions. And this is kind of one of the more clear-cut cases. But we're also starting to see-- and I think we're going to see more of this-- for example, OK. So, that's an extreme case. What about

**AUDIENCE:** People who plead insanity after doing something really stupid?

**ABBY NOYCE:** Insanity plea is an interesting thing. Legally it's got to be based on a professional's judgment that the person did not know that what they were doing was wrong at the time they did it. So, the classic good case of this is, for example, somebody who is severely mentally retarded and doesn't know that it's wrong to take a candy bar, for an example of this.

**AUDIENCE:** But what if somebody's a sociopath, who doesn't realize that it's wrong to kill people?

**ABBY NOYCE:** So, that's where you don't send this person to the death penalty. You send them to involuntary committal. You move them-- so, the state has the right to-- if somebody is a danger to others, the state has the right to get a judge to say, this person can be locked up. This person can be confined to a treatment center, until they are no longer such a danger, which may be indefinite.

So, what about something that's kind of like the brain tumor issue. So, we know that, for example, in a lot of ways, experiences in early childhood really strongly influence a lot of how your brain develops. For example, we know that kids who are abused as little kids tend to grow up to be more aggressive, to have a hard time forming relationships, to have, in general, have more issues, all around, than kids who grew up in stable, healthy families.

So, what happens when you have some young adult who does something violent, who hurt somebody and the case that their attorney makes is he couldn't help it. His brain was wired to do this by his past. Where people start saying that you lose-- that because of how our brains are shaped by experience that happens to them, and then our future behavior is defined by these experiences, how much control does one really have over what one does? It starts getting kind of iffy. And it's kind of dark to think about, in a lot of ways.

**AUDIENCE:** But then, wouldn't that mean that no one's responsible for anything that they do?

**ABBY NOYCE:** Yes.

**AUDIENCE:** They can go around killing everybody, knowing--

**ABBY NOYCE:** This is the slippery slope argument for this, is that it's definitely possible to start out thinking about this and come to the conclusion that there is no such thing as personal responsibility, that you cannot be held responsible for anything you do because it's all just your brain doing it and it's all defined by personal experience, previous experience. It's all mechanical, that there isn't a choice option involved.

**AUDIENCE:** That's why I think that, legally, they should not take into account-- they should assume that everyone does have free will. Because otherwise, it wouldn't work. Nothing would actually work.

**ABBY NOYCE:** Yes, this is an excellent point. That your assumption of free will-- not just for the legal system, but for, like, getting out of bed in the morning, dude. You've got to believe that you have free will or it all falls apart. I don't know.

**AUDIENCE:** --slave to your brain.

**AUDIENCE:** --robots.

**ABBY NOYCE:** So, what's the difference between you and your brain?

**AUDIENCE:** I don't know.

**ABBY NOYCE:** Is there a difference between you and your brain?

**AUDIENCE:** It's kind of squishy.

**ABBY NOYCE:** It's a good thing it's inside that nice solid skull. Is there a difference? Part of the big question that comes down here-- is there a difference between your self and your brain? Is that a distinction that we can make validly?