

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.

PROFESSOR: OK, good afternoon. So today we're going to talk about learning. So what can you do now that you couldn't do as an infant, very simply?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Talk, that's a good one. Language, that's pretty good, right? Anything else that you couldn't do, yeah?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah, we don't actually know what the internal mental life of an infant's like. But it doesn't seem like they remember many very specific things about their lives. You do. But you know it's a lot more than, you know a ton of stuff, a lot of facts, right, things about the world, the values you hold, about your family, your culture, the country or countries you've grown up in. You just know a ton compared to an infant.

And there's only two ways that you know this. It's either in your genes when you're born or you learned it through life. And so we're going to talk today about how scientific psychology tries to approach learning. And in a remarkable thing about learning is we often think, well, we're learning lessons or whatever. We're learning material.

But the important part of learning the way we think about it is that it allows us to predict the future on the basis of the past, right? We want to learn lessons in life to know what's desirable, what's dangerous, what's a smart way to do things so we're more effective. We choose what we want to do better in the future to learn from the past and predict the future.

It imbues our word with meaning. For an infant, not much means much, right? They

just look around. Who knows what anything means. Words, gestures, facial expressions, there's not much meaning of the world out there. For you, the world is full of meaning, history, social interactions, everything that you're thinking about. So how do we acquire the meanings of things in the world that matter to us?

And it's also had an interesting kind of a give and take with how we learn about learning in scientific psychology and sort of historically. I'll just touch a tiny bit on that. And we'll talk about three aspects of learning today: classical conditioning, operant conditioning, and then various ways in which there are limits to conditioning. Thank you, Todd.

So let's start with sort of one of the most famous figures in all of conditioning or learning, Ivan Pavlov, who actually won a Nobel Prize, not for his studies of learning, but for his sort of groundbreaking studies for the reflexes of digestion. He worked out a lot of the fundamentals of how, in mammals, your stomach breaks down food and that how food in the mouth provokes specific salivation that begins the process by which food is decomposed and then digested in your stomach, the fundamentals of really how we live, right, eat to live. And he was very interested in salivation reflexes, what is it that drives that first bit of breaking down food in your mouth with saliva?

And Pavlov was a pretty intense researcher, as you might be used to at MIT. There's a famous story where a graduate student came somewhat late to the lab and said, but Professor, there's a revolution going on with shooting in the streets. This is in Russia. And he said, what difference does it make when you've got work to do in the laboratory? Next time there's a revolution, get up earlier; so anyway, not an easygoing supervisor.

He would do experiments a little rough in dogs, like cut the esophagus, the path by which food goes from the mouth to the stomach, to understand the role of the esophagus. And he found that, even in those animals, when he placed food in the dog's mouth, he expected, as a fundamental biologist studying how digestion works, that nothing would happen.

But he noticed not only would the dog salivate, even though the food will never make it to the stomach, but furthermore the stomach juices, the gastric juices that break down food, also were released within the stomach even though the food would never get there.

And then he got interested in how it is that these things are driven if they're not driven simply as a direct response to the arrival of the food. How is it that the mere sight of the food or the sight of the person bringing the food would drive the salivation and what he called psychic secretions?

Because he was a biologist. So he expected, basically, you put the food in the mouth. There's some chemistry between the food and salivation. And here he's finding, just seeing food, just seeing the person who would bring the food, was driving these fundamental biological functions.

So there's a kind of experiment that he would set up where the dog would be there. And he would be measuring the amount of salivation in some way or another. And in these kinds of experiments-- and I'll go through this language a couple times-- in conditioning they've come up with a sort of vocabulary to describe the pieces of things that go into this kind of learning.

So there's an unconditioned stimulus, which is food. I mean you start with that. You want to eat food. And there's an unconditioned response, which is you salivate when you see the food or if it's in your mouth. That's the beginning.

Then you're going to add what's initially a meaningless stimulus. In this case, it could be anything. But in this case, we'll say it's a bell. You ring a bell before the food comes.

And you're going to get to a conditioned response. The animal will say, well, the bell predicts the food. And I'll salivate to the bell itself. Even though I'll never eat the bell, I learned that it predicts the food is coming.

And so you've created a new association. You learned something new about the

world, that a bell signals something that drives salivation. So the two things are sort of occurring next to each other in time or what Aristotle talked about as the law of contiguity. The animal learns that the bell predicts the food coming.

And so here's the basic idea. At first, the tone does not drive salivation, right? There's no reason to salivate when you hear tone by itself. But as the animal discovers that the tone predicts the food, the tone predicts the food, the tone predicts the food, then all of a sudden the animal starts to learn that simply the tone alone will drive the salivation.

Now this stimulus is conditioned. And you have a conditioned response that wasn't there to start with. It's very simple, right? You learn the two are associated. And now you begin to have a biological response to the tone alone.

And typically what happens is here's a learning curve. Here's how often perhaps you salivate. And across trials you learn more and more of that. And you finally plateau at some level of performance.

So a neat thing about Pavlov is, although there weren't a lot of movies at his time, there is a little bit of film of him and his actual dogs. So here's Pavlov's actual dogs.

So we said conditioning is going to do two magical things, right? It's going to let us predict the future on the basis of the past. And it is a very simple experiment. Now the bell predicts the food coming up and the salivation.

And it imbues the world with meaning. In a sense, the bell now means, here's dinner, right? The bell is not just a bell. It's something with meaning and significance.

These are another example of classical conditioning. And it's got some sophisticated properties. You could go, well, balloons popping water in your face, food for dinner, that's pretty basic stuff. But here are some sophisticated properties of these as learning mechanisms.

So one was already talked about, extinction, that, if the bell continues to ring and no

food arrives, the learning gets weakened, weakened, and disappears. So you can learn something. And you can unlearn something because it's no longer an effective predictor of the future.

Generalization is kind of cued also. So imagine the frequency of the stimulus is this, just the specific frequency of the tone. Well, if the tone is a little bit different, you might still think food is coming, right? So it's not that you learn the precise thing. You'll still respond if it's a pretty similar tone. And as the tone gets more and more different, the response gets weaker and weaker, right?

So that's smart generalization. I stay pretty close to the tone. I go for it. I get further away, eh, not so sure. It's not all or none. It's reasonably graded and sensitive to the specifics of the situation.

There's also one fascinating thing, which is your acquired conditioned response is extinguished because you're only hearing the bell and there's no more food coming. But when you come back in, you often have a kind of preserved recovery function of that CS. It's like it's lurking in you somewhere.

It's like I have a mental note that this bell might mean food, even though it was extinguished. It's still pretty interesting stimulus. So you'd have some easily spontaneous recovery of these kinds of learned conditioning.

It can also be sophisticated in other ways. And let me give you a couple of examples. You can learn to discriminate between two things. We've only talked about one thing so far, bell, food. But what if you have a black patch, for example, that predicts food coming, and a gray one that predicts it's not coming, you can distinguish between those two.

You can get second order conditioning. And let me show you that. The dog learns that the bell predicts food. Now, there's no food presented for a while. But there's a black patch presented that you see presented with the tone. Because you know the tone predicts food, the animal figures, well, maybe this is a pretty good hint that food is coming.

And after a while, it begins to salivate to this one, even though it was never directly paired with the food itself. So it's learning. It's a transitive inference that it's making. If the bell predicts it and the square goes with the bell, at some point I'll start salivating for the square because that sounds like a pretty good candidate for predicting the future.

So that's pretty sophisticated, too. You don't even have food in this part. But it's even more sophisticated than that.

So look at these two things. And you can think about, I'm going to let you think about this for a moment. Here are two groups of animals. Or it could be people. It would work just the same way. So here's the conditioned stimulus and the unconditioned stimulus, for example, a bell and the food, the bell and the food, the bell and the food.

One group only gets bell, food, bell, food kinds of pairings. The other group gets the same pairings. But in between, they'll have a random food reward, two random bells, a random food, random bell. Now, the pairings are identical. That's what's shown in orange. This group will learn better and faster, more strongly, than this group. Even though they've had the identical number of pairings.

Why will this group learn more slowly? Why will, even though they have an identical number of times that they got the tone and the food, the tone and the food-- Yeah?

AUDIENCE: Is there a little bit of extinction between--

PROFESSOR: It's an idea that is an extinction between-- I think that's a good idea. So that's a very good answer, the extinction between-- because, when you get this CS alone, you're fighting the learning. But it's telling you the system is smart, right?

Because what it's kind of doing is saying, look, eh, CS alone, not always a reliable predictor. So I'll learn it. But I'm not going to believe it so strongly, right? It's basically because you're sticking extinction in between those, right? You're also having some trials with the food that aren't predicted either, right?

So it's like the animal or you are learning what's a really excellent signal. And if there's other stuff in there that gives the signal not the same predictive power for the future you learn more slowly unless certainly that relationship. OK? So that's pretty smart, too.

Here's another way in which it's smart. And I'll tell you the phenomenon. And you think about why you think learning might work this way in animals and people. So pretend there's a tone that predicts food just like we've been talking about all the way through. Food stops.

And now you get a tone and a light, a tone and a light, a tone and a light. The tone predicts it's food. You're getting this part, again, without food being presented. And now you get the light alone. OK? Or everything's the same. But there isn't the initial learning with the tone.

And then in the end you're tested for the lights. And what happens is there's less conditioning to the light here than here, even though you've seen a light an exact even number of times. Why is the conditioning to the light weaker here? Why is it weaker here than here? This is an information problem for the brain to solve about the truth of the world.

Well, our interpretation is this. Tone, food, tone, food, you go, good, I got it. Easy, right? Now you get tone and light together for a while. And you go, OK, I know the tone is really important. The light, I don't know. The tone was perfectly good. The light seems redundant, informationally redundant.

So now when it's a light alone, you don't have so much conditioning. Because you didn't care about the light so much. The food had all the information you needed. The light was informationally redundant, right? It was superfluous.

Here you only got trained with the tone and the lights. So you never made a bet on what's the critical thing. They both predicted evenly. So when the light comes by itself, that was always something that seemed like a pretty good signal equally good to the tone.

So that's very sophisticated that the animal learning it or the human learning it is deciding what's the information that's diagnostic and what's the information that's redundant. And it's conditioning across that, right? So it's not just very simple reflexes. It's sophisticated learning.

So now I'm going to switch gears. So we're going to have a demonstration that involves the opposite of fearful things like spraying water and popping balloons. But I can't tell you too much about this. But I can tell you it'll be pretty pleasant.

But I need somebody who's willing to step out of the room for a moment to do this. OK, thanks. Yeah, yeah, you can safely take off your jacket now. And what's your name?

AUDIENCE: I'm Sam.

PROFESSOR: Sam, OK. Try not to listen if possible, OK?

AUDIENCE: Yeah.

PROFESSOR: So you guys are going to help me perform instrumental conditioning. And then we'll discuss what that is. Our job is to have Sam come over here and pick this up. OK?

Your job, as Sam comes in for every action he takes, a step this way, you played this game as a kid, or an arm going in the right direction, applaud if he's doing something that gets him closer to the goal of picking up the brain and withhold your applause if he's doing something that's taking him away from that. Is that OK? All right, OK.

So we're going to discuss what instrumental conditioning is. And then we're going to slightly think philosophically about whether this is how you learn to do everything. You get to choose who the applause comes from, your parents, your family, your friends, medical schools and law schools and graduate schools. You get to choose who the applause comes from, about how much of our lives is responding to the applause that we care about.

So we'll start more simply, less philosophically. Let's think about this. The examples we talked so far in classical conditioning all build on what feels like a basic reflex, salivating for food, flinching for water, right? So if we ask somebody, why do you work hard, there's not to be very good explanation in terms of those initial reflexes, right?

You're not flinching. You're not looking for food or water. And yet much of our life is involved in endeavors that we learn where there's no primary reflex that we can see in the story. And so operant or instrumental conditioning, those two words are interchangeable, is the way that psychologists have tried to understand those forms of learning.

And it began something like this. And this is from Thorndike. He put a cat into a puzzle box. I'll show you what the box looks like or a sketch of it. The cat had to unlatch a door by pulling a latch.

So here's, the cat's in a box like this. To get out, which cat wants to get out, he or she has to push this, which will push the latch. Now, that's not obvious to a cat. It wouldn't be necessarily obvious to a person but definitely not obvious to a cat.

And the animal learned it by trial and error sort of like Sam did, right? And where is the unconditioned stimulus? There's no salivation or any reward like that. So Thorndike proposed that people and animals learn things by the consequence of their response.

And the Law of Effect says the consequence of a response determines whether it is strengthened or weakened. You can be rewarded. That will strengthen a response. You can have no reward. That will weaken it. You can be punished. That will greatly weaken the response.

So here's the cat. Here's an example of a cat. At first, it takes a long time for the cat to get out of there as it tries different things. But over repeated trials, the cat can get out really, really fast.

And how Thorndike, by observation and thinking, thought about it was at the

beginning the animal wants to get out. It has no idea how it gets out. So it does everything it can. It scratches at the bars. It pushes at the ceiling. It digs through the floor. It howls. And at some moment, it presses a lever. And that's good, right?

So here's different things the animal starts with thinking might be good ideas. But because it discovers with repeated efforts, repeated trials, that pressing the lever works, these get weakened. This gets strengthened. So the consequence of a response determines whether it's strengthened or weakened. And now there's no salivating or reflex or anything like that. You can learn then almost anything.

So this was picked up a lot in the US. John Watson is a famous name. You'll see kind of a haunting film of him performing this kinds of stuff with infants in a kind of a way that we would never allow ethically now. It's not horrible. But it's pretty bad. But it's a famous thing. And you'll just get a feeling of the power of this.

Behaviorism, that if you study behavior, instead of going for things like unobservable thoughts, that you would just do observable actions. And you wouldn't make inferences about the mind that you can't measure. You make a stimulus. You measure the response. And that's it. No fundamental differences between animals and humans, and you make laws that describe the relationships between stimulating the environment and physical behaviors that you can observe or responses out in the world that an organism has.

So I'll come back to this a little bit. But one of the ways he wanted to show that anybody could learn anything, because right now, don't forget, the way we set this up, anybody can learn anything. You can applaud anything, right? So you can teach anybody anything by this perspective.

And he wanted to show that he could teach an infant initially, whose name was Little Albert. That was his cover name. Nobody actually knows his real name or what happened to him. Well, nothing horrible happened to this person we know about.

And now you would never do an experiment like this. You have this in your notes. This is a long version of the movies taken of this actual sort of experiment that we

wouldn't do now. And I'll show you a short version in a couple minutes where he got an infant to become afraid of a rabbit that the infant previously was not afraid of.

And then the other huge name in this is B.F. Skinner who worked at Harvard, consequence of a response, again, creates the responses that we make in the world. And his sense response is not just the laboratory one but things we do in the world with each other and to each other.

And so the classical conditioning, the CS elicits the CR. The conditioned stimulus, the bell, elicits the conditioned response, salivation. But in instrumental conditioning CRs are omitted. And you can do really anything that's in the range of what an organism can do.

And then he could show kind of impressively how you might imagine a chain of these things would lead to complex behaviors. So he would teach pigeons and other animals to first to click to get a pellet, then where to click it, then to where to face before you click it. And he would just keep adding little pieces, shape the pieces up from the first step to the second step, until he got pretty complex behaviors.

OK, so now I'm going to show a slightly longer clip that has in it these kinds of experiments and Little Albert. It's kind of cool because in the sense that we can actually see films of these sort of historical things, which often we don't get to do in various science fields.

So let's talk about that for just one moment. So we can talk about primary reinforcers like food, thirst, or pain, things that go with life, right, pain, not being injured, food, and thirst. We can talk about secondary reinforcers that we know a lot about, money, attention of other people, praise from other people, admissions to places you want to be admitted to, promotions when you want to get them. There can be positive or negative rewards. And I'm going to talk about one thing.

But just for a moment the behaviorists said anybody can learn anything, right? We'll come back to that. So that's been a big debate whether anything can be equally learned with any other thing.

And the other thing that's striking is B.F. Skinner particularly said that we have an illusion of free will, that all we are, are learning these things. And we might as well get over our illusion of free will.

Scientifically, it's very hard for even the best scientific psychologists to tell you whether free will exists or not. That's not something that's easy for us to begin to imagine how to measure. These kinds of learning things are very impressive in many ways. And we'll talk about some of that.

So let me talk to you about one way in which, for example, casinos get people to come back often and other things in life, without giving you money every time. Because then they wouldn't make money, right? So your book talks about a number of examples of this. But I'll just pick one, partial reinforcement.

So this means when do you get, in these animals, when do you get a reward? Every time or only 1/3 of the time? So of course, you learn faster if you get a reward every time than if it's 1/3 of the time. You're not quite as impressed and excited.

But here's the interesting thing. If you stop giving the reward, these animals will go on for quite a while before they extinguish. These animals will extinguish right away and stop responding. Why will these animals keep behaving for a time, even though you've entirely stopped the rewards? Yeah?

AUDIENCE: Since they only got it 30% of the time anyway, it seems plausible to them that it would go a couple more times without them getting it.

PROFESSOR: Yes, since they only got, the answer was, which was right, they only got it 30% of the time, so the first one or two nonpayments, right, they're going to go, well, it's coming up next, OK, in a couple more, OK, I'm coming up soon until they finally decide that the whole thing is over, right? So that's very interesting, the different reinforcement schedules, how they relate.

You might think, well, you always learn better if you reward more. But in some cases, if you want to get people to keep doing stuff without direct reward, partial reinforcement is actually more powerful.

Now we're going to move into dog experiments and then directly to sort of thinking about people and our own lives. So this is a famous, famous experiment at the University of Pennsylvania from Seligman. In this case, he used dogs. It's been done with many species in many circumstances.

And he put them in a hammock. So there's two dogs on two sides of a thing that they walked around. And they don't see each other.

The dog, A, would periodically get a shock. And it could stop the shock by pushing a panel near its nose. As soon as the shock came, it could turn its nose and stop the shock.

Dog B on the other side of the divider would get exactly the same shock. It would stop at exactly the moment that the first dog stopped it. But dog B had no direct control, equal number of shocks, equal duration of objective pain, if you want to call it that way. One animal has the capacity to stop it. The other does not. So that's the first part of the experiment.

And the second part of the experiment they're put into a shuttle box. It's just a box with a big divider in the middle. And they're just hanging out there. And then they hear a tone.

At first, it means nothing to them. And then comes a shock on the grid on the floor. And at first, they don't know what to do. But then after a while, they get smart like the cat. And they jump over the divider to the other side. And it's all OK.

So you can imagine that when they hear the tone pretty quickly that what do they learn to do? Jump over the divider and stop the pain. And here's the remarkable thing. For the animals who had the experience that pushing the panel stopped the pain, they learned that pretty quickly. For the animals who could not stop the pain initially, they don't learn how to escape when they could.

It's as if the animals in group B, what they learned about life, it's painful. And you can't do anything about it. The animals in group A learned pain happens. But I can

do something about it.

So here's the animal. At first, it doesn't know what to do. It's getting the shock. But with trials, the animal learns to jump and jump so quickly that it doesn't have almost any problem. Here's a tone, boom. It's out of there. It learns to avoid the pain.

But what happens to the animals who either had the escapable shock or the ones? The ones who could poke their nose, they also learned this. The ones who could jump over, there's nothing preventing them from learning that. They just don't escape.

It's what Seligman called learned helplessness, that the animals learn, from the first experience, their conclusion was pain happens. And I can't do anything about it. And then when they get in a second situation where they can do a lot about it, they continue to feel helpless and not do something about it.

It's like a deep lesson about life, like when do you think you can make your situation better and when do you think there's nothing I can do about it? It's just bad, bad, bad.

So there's ideas about that some of this might be related to people who struggle with depression but also other things in life. So here's the idea from Seligman as an explanation of how we have resilience. The idea is that we all get in sometimes in small doses, sometimes in tragically large doses, difficult things into our lives.

How do we explain those difficulties that have happened to us, the setbacks we have? Do we think they're internal to us? I'm the kind of person who has bad luck. I'm the kind of person that these things happen to. Or is it external? Somebody gave me a test. What can I do about that? But next week the test is over. And life will be better.

Is it global or specific? Does it just seem miserable all over? Or is it just one specific squirt in the face that this faculty member gave me? Is it stable or unstable? Are bad things just seeming to happen all the time? Or is it one occasion?

And you can imagine that people who believe that bad things are internal, global, and stable will learn to be helpless. Everything seems dark and hopeless. People who believe that things are external, specific, and unstable think, the next time I can do better. I can bounce back from this and do better.

And it's not only in terms of the sort of spirit of people and trying to get people to bounce back from setbacks. It's even things like job applications. Seligman did the following bet. I don't even know what corresponds to it before.

But there used to be an era when people would go knock on doors and try to sell you insurance. We don't do that anymore, right? Because if somebody comes to our door, knocks on our door, we're calling the police. But when things were more calm, you'd open the door and talk.

Or they would give you cold calls, which you don't get much more anymore these days. You guys, have you ever even heard that phrase? You'd work at some insurance company. And you'd go, here's some random telephone numbers. You call them up. And you try to sell them insurance. Usually, you get abused on the phone. Or somebody slams the doors.

But these salespeople, they make commissions. And they keep going. And he said, let me get my test. He gave a questionnaire whether people interpret the world one way or the other way, the resilient way or the learned helplessness way. And you give your usual interview way for picking sales people. And let me see who picks a better salesperson.

And his questionnaire picked a better salesperson. Because if you have a sales job where people often kick you out and hang up on you, you better be pretty external, specific, and unstable in your perception on things, right?

Because when you have 100 people who hang up on you or close the door on you for one sale, you need to have a thick resilient skin. So there's lots of ways in which I think this speaks to how do we respond to the setbacks in our lives.

So let me talk about some limits to conditioning. And some of these that I think

touch also on what are the values by which we lead our lives. It's not going to be telling you that. But it'll make you think about some of the mechanics of what matters to you.

So let me slide back for a moment. So the behaviorists said everybody can learn everything. Everybody can learn everything. In that sense, it was seen as a very democratic and egalitarian perspective. Everybody can learn everything. It's just how you set up the contingencies in the environment.

And then Garcia did a famous experiment where he had rats get shocks or receive lithium chloride, which made them nauseous. And they either had a bright noisy stimulus. Or they tasted some sweet water.

And even though these things were arbitrarily combined, the rats learned, who received shocked, learn to fear the bright noise, the bright light and the noise. And the rats who got the lithium chloride that made them nauseous learned to fear the sweet water. Do you see why?

When we get nauseous, do we think like, boy, it was noisy last night? Or do you think, when you wake up in the morning and you feel nauseous, what's the first thing you think?

PROFESSOR: Yeah, something I ate, right, we're prepared for internal nausea to think about food as a risk. And if it's a shock, we're prepared to think of something out there. So these animals instinctively, not on the basis of the training, instinctively said, if I feel a shock, it's probably related to this bright noisy stimulus that warns me. If I feel nauseous, it's related to the sugar water.

And they had these opposite patterns of learning for the stimuli as if these animals were prepared to fear some things and not others. So it wasn't that everything could be matched with everything. And in fact, with humans, they've done things like show you pictures either of snakes and spiders or flowers and mushrooms. And then those predicted a shock, a small shock to humans.

And they measured your galvanic skin response and autonomic response. When you think you're going to get a shock, you start to get a little sweaty. And what they found was-- and this will not shock you-- they had better conditioning, you more quickly and completely learned that a snake or a spider predicts a shock than a flower or a mushroom. Even though, objectively, they were just as predictive.

There's something about snakes and spiders that gets us ready for scary stuff. And flowers and mushrooms, they really have to punish us before we think they're evil signals so, again, a human example of preparedness. We don't learn everything equally. We are prepared to learn some things and not others.

Same thing with Little Albert, the child that you saw, that kind of a film has a funny status. And we'll talk about a few more things this semester where that experiment, why should that experiment not be done? Yeah?

AUDIENCE: Nobody knows what happened to Little Albert.

PROFESSOR: Yeah, well, you would never terrorize a child, right, to scare the child. So weirdly enough, though, these studies remain in the field because since we can't do them again. Because we know they're not right to do. We try to learn lessons from them so that we learn something.

The rat worked. He got scared of a rat. And then he got scared of a bunny. He got scared of a cat, anything animal, anything furry. But if they gave him a wooden block, that didn't work. It's not everything. It has to be prepared from that.

Here's another one that also challenges simple ideas of do we always learn simply to get a reward. What kind of learning might exist that's rewardless? So there's three groups of rats in a goal maze.

There's food rewards every day for one group. They're the happy group. There's no rewards for a grumbling group. There's a third group that gets no reward for 10 days. And then it gets reward.

So reward all the time, reward never, no reward for 10 days and then you start to

get rewards all the time, and here's what happens to the learning. If you look at this green line, that's the food that gets rewarded all the time. So they make fewer errors so they can get to the food faster. If you look at the group that's never rewarded, this yellow line, it's pretty steady. It's like I'm hanging out, no reason to go anywhere. There's not much going on.

Look at the orange group. For 10 days, up to here, they look like the yellow line, right, not getting rewarded. I'm just going to saunter around. There's nothing interesting to learn here. But boom that first day that it's rewarded, but just one day of learning, they're all caught up. They're all caught up.

So people call this latent learning, that they were really learning a lot about their environment anyway, even though it was of no use at the moment. And then the moment it became valuable information, boom, they're all caught up. It's as if they learned all this information. And it was ready to go.

Latent learning, the information is stored up just because they were in the environment, even though it was unrewarded. And as we go around, we often pick up stuff, even though there's no direct reward. And sometimes it turns out to be useful.

People also love, humans, including infants, contingency, the idea that we feel like we have control over our world. So here's a two-month-old infant, very young infant, in a little experiment where they put the infant in a crib. And above is a very colorful colored mobile.

And in one version, when the infant moves his or her head, there's a switch in the pillow. The mobile moves. And the infants are smiling and cooing. They love it. This is awesome. They move their head. The thing moves, pretty impressive for a two-month. All they can do is wiggle their head, right? They're not walking or anything.

They put another group of infants in who have no switch. And just periodically, they turn on the mobile. There's no smiling and no cooing. And they equate the number of mobile turnings. Why is that?

Well, if it's just because you like the mobile turning, then you should have both groups be equaling cooing and smiling or disinterested. But there's something fundamentally rewarding for humans, even at the age of two, to feel like they have control over their environment and their pleasures.

And that's not explained by conditioning. Nobody's getting anything more. They're both seeing the equal number of interesting turns. But one group feels like they're controlling their environment.

Here's another one that I'm going to show you in an animal experiment where you can control things. But you can imagine in your own life. In my life, I see this all the time. It's nice when something pleasant happens in your life and your situation gets better in some way, right? You get a better iPad, right?

But how hard is it emotionally when something gets worse? They take all your iPads away or more serious things. When things get worse, often we feel pretty miserable. And here's a neat example of that.

So these are rats in an experiment so everything's controlled. So here's one group of rats. This is how quickly they're running. So it's good to be high. One group of rats is constantly getting one pellet a day. So they learn what to do. And they're sort of hanging out here doing well.

A second group of rats is the luckiest group of all. They get eight pellets. And they learn more. But they flatten out. So you get more reward. You drive the behavior more. That's not unexpected.

The interesting thing are the two contrast groups. The contrast groups start off with one or eight. And then, after a couple days, they're flipped around. So some started with one. Then they got eight pellets. Others started with eight. Then they got one. Does that make sense?

So here's the contrast group of eight. They're learning one. They go eight, woo. They're happy. They're getting eight. Life got good.

Look at the group for whom life got worse. I'm getting eight. I'm getting eight. I'm getting eight. I'm getting one. I'm getting one. I'm getting one. It really stinks.

They're actually worse than the groups that got one all along, right? This is the lowest line of all. They're protesting. This stinks. I was getting eight. Now I'm getting one. What a rip off. What kind of world do I live in? So the learning is protesting that.

Now, by purely conditioning one, why would this group be worse? They're still getting a pellet. But because they're interpreting it as a world that's going downhill, as a worsening situation, they're less prone to learn it. They're sort of protesting in their learning, things having gotten worse in the world.

A really interesting theme along the line of this reward, an unbelievable one that has gotten a lot of attention, is delayed gratification. If you think about it, and the analysis is not purely scientific but is a theme, delayed gratification is a spectacular part of your life, right, especially if you go through higher education.

Or maybe everything is like this, right? Do a good job in primary thing and you'll get to a good middle school. Do a good job in middle school, we'll get to a good high school. Do a good job in high school, you get to MIT.

Do a good job at MIT, you go to med school or law school or become a teacher, whatever you want to do. A good job at that and you'll get to this thing. Do a good job of that, finally, thank you very much, you're 95. And you're done. I mean it's kind of like you guys are waiting a long time between the efforts you're doing and what you might consider a palpable reward.

So here's a fantastically interesting study from Walter Mischel done at Stanford some years ago. These children were four to five years old. There were about 600 that were studied. They were mostly the children of faculty and graduate students at Stanford.

And they would come into a room. And the experiment, slightly varied, was basically like this. They would sit in front of either a cookie or marshmallow. They used different rewards. But imagine it's a marshmallow. In the movie I'll show you, it's

Oreo cookies.

And they were told this. And they tried to make this, Walter Mischel who did the experiment describes how he played with kids to try to get the marshmallow look really awesome. And the kid would sit there, the four or five-year-old. And they were told, I'm going to leave the room now. And there's a bell here.

If you really want to eat that marshmallow, ring the bell and just eat it. But if you wait a while-- they didn't tell them it was 15 minutes. It was 15 minutes. I'll come back. And I'll give you two. Delayed gratification, I can have one marshmallow now. Or if I hang in there for 15 minutes and don't touch it, I can get two.

Let me tell you the results. And then let me show a film of this. When I was looking on YouTube and the internet and stuff, there's a lot of posed movies out there. This is the only one I found from Walter Mischel himself that's a real movie of real kids really doing this. I don't know why all the ones are posed out there.

But here's something astounding. So these are all pretty similar children, children of faculty. Children varied how long they could wait. And if they took, for example, the children who didn't wait very long. They could only wait 30 seconds until they ate the cookie. And they compared them to the children who would wait for the full 15 minutes to eat the cookie. These are four and five-year-olds.

When they looked at them at 18 and their SATs, the children who would wait 15 minutes scored over 200 points higher, which you know is a lot. That was in the days we just had two SAT tests. So it was 1,600.

And by many, many measures in terms of educational attainment, in terms of health, and many measures, the longer you waited at age four or five, the healthier you were, the better SAT scores you were, the further you went in school, astounding correlations with real life stuff. And people are stunned by this.

Partly, it's again, who knows what goes into that at four or five. How much is the influence of parents? How much is the influence of cultures and genes? Who knows? But that fateful moment when the child goes for the marshmallow or not

correlates with incredibly important outcomes in the future. So let's see. Here we go.

And you know, right, every college I remember personally, it was always a battle, like am I going to have fun or study? Am I going to have fun or study? There's so many chances to have fun. And there's so many pressures to study, right? Well, studying can be fun. Learning is fun.

How about when rewards harm? Again, these are things that are not strictly predicted by conditioning. I have three or four more minutes. So rats love to run, if you had a hamster or anything like that, they love running wheels. What if you give them, every time they run, a food reward?

Then it turns out rats no longer run for fun. They only run when the food is given. If you start to extinguish, they stop running. So what had been a natural thing that seemed like a whole lot of fun, if it became the price you pay for a reward, all a sudden becomes a chore.

A widely cited experiment for Mark Lepper, he took preschoolers, so four-year-old children, who love to draw. Lots of preschoolers love to draw. Then they started to give them, every time they drew, a gold star. Every time they got a reward.

What happened when they stopped given the gold stars? The kids stopped drawing. This is incredibly often cited in incredibly interesting debates about whether children in impoverished schools should be given cash rewards for studying hard. There's lots of debates about whether that's a good thing or not a good thing, an effective thing or not an effective thing.

This study is often cited for arguing, if you make something that you should love, like learning, purely rewarded, the minute the reward stops, it becomes a complete chore. But we don't know how that applies. Because we might read about things like baseball player salaries. And it seems like there's a pretty good relationship between performance and salary.

So it's not simply that way in the world. We know that often rewards are associated

with high performance and promote high performance. It's pretty complicated.

The last slide is about a fantastic debate that occurred here in many ways between B.F. Skinner at Harvard and Noam Chomsky at MIT, which is, in what sense is language learned by conditioning? So here's an example against it.

At one month, there's a switch inside a rubber nipple hooked to a tape recorder. And when the baby sucked, the tape plays. This allows the baby to perform, to behave. And they would play syllables from different languages. And they may or may not be in their own languages.

And by four days, babies preferred the languages that they are exposed to. And one of the questions is, people said, well, how do we get kids to learn language from a behaviorist perspective? So one of the things that they studied children, and they'll often, young children will say wrong things like mama isn't a boy, he a girl. That's not good grammar, right?

And so there was an intuition that parents would then go, that's not good grammar. And sometimes parents do that. But when they did studies where they recorded long sessions with parents doing this, parents almost never corrected. It was very rare. You may feel different from your own record.

But statistically, it's very rare to collect. They often say that's right or something like that. So how's the child picking up language if nobody's giving them rewards in terms of whether they're right or wrong?

The idea is that we can say an infinite number of sentences. Well, how can we possibly be rewarded or conditioned for an infinite variety of sentences? How could that work?

And finally, the phenomenon of overgeneralizing, you may know that children will overgeneralize things. Because we usually end E-D to make a verb past tense, right? But some verbs don't work that way. So you would say, my teacher held the rabbit. Children often make these overgeneralized rules. They say, my teacher holded the rabbit.

So obviously, they're learning a rule and then applying it to an instance that they never hear. Because almost no adult says that. So they must be producing that, not on the basis of conditioning, but Chomsky argued on the basis of sort of genetically prepotent ways in which we learn language, rules that we develop and figure out language in the world.

I will say that, for those of you happen to be interested in this topic, there's been a striking pushback on that in what people call statistical learning. So as much as this seemed like a done deal 10 years ago, it turns out the truth might be somewhere in the middle with some genetic prepotentization but also a lot of learning from your environment both.

So you can think, as you leave here, what are the applause that you listen to? Where's your free will and your choice of what goals you pursue and what matters to you? And how does the world attempt to teach you things that you learn about what you want to do and why you want to do them?