PROFESSOR: Anyway, if you're sitting there saying, what are they talking about weird thing in the syllabus, that suggests you didn't read the syllabus, which is why we put the weird line in the syllabus about-- what is it? What a warrior your mother would have been? That sound faintly familiar?

Well, it should sound familiar from reading the syllabus, right? Nobody's got a clue where it's from, right? Except for Rachel. I know Rachel knows because I told her today.

And obviously, I cleverly arranged for it not to be something that could be Googled. Because several years ago, when Google appeared, I had a line in there to make sure people were reading it. And within tens of handing out the syllabus, half the class said, oh yeah, that's from like Richard III or something. So this one couldn't be Googled.

It is, in fact, from a book by John Steinbeck. You probably all read some Steinbeck in high school. Those of you who are fans of Arthurian literature might want to read his rewrites of the King Arthur or chunks of the King Arthur story. And when you do, you will discover that that line about thinking about what kind of a warrior your mother would have been is in there. Great book.

Any other questions of similar moment that I should be answering at the present time? No. OK, well, what I'm going to talk about today is more of the detail, if you like, of the-- last time I had up on the board this business about, you're a slave to the environment, you're a slave to your brain.

Well, this is the details of the way in which you are a slave to your environment. Now, typically, that is-- the heading of-- the title of the lecture is something like learning, the chapter is learning. You're here thinking that what you came to do today is to learn, and that's not what we're talking about.

The sort of learning that you are doing here is material that really gets covered in memory. I present material to you. You store it in memory somewhere. At the appropriate time, you retrieve it from memory. We'll talk about that in a few sessions down the line.

What I'm talking about today is a particular form of learning, which you can think of as association learning, a very reflexive form of learning, a form of learning that is phylogenetically very old. It shows up in all sorts of-- essentially, in any beast you care to try it out in.

Indeed, if you were to go to-- those of you who are budding neuroscience sorts, if you were to go to Eric Kandel's website and look at his Nobel Prize lecture-- because he won the Nobel Prize. You don't usually post your Nobel Prize lecture until you win it because it's a little arrogant otherwise.

He won the Nobel Prize for working out the detail of simple association learning in a beastie known as-- initially, his work was done in a beastie known as aplysia or aplysia, which is a sea slug. It looks like a dill pickle and has just about as much personality and brain. But it is capable of this form of learning. Gorgeous.

It's a gorgeous lecture. Nice art on the website. And if you're going to give a Nobel Prize lecture, you don't want to waste it. So I heard him give a version of this at a different meeting in Australia. And it's a good lecture.

All right, what kind of learning is this that we're talking about? How many people here like strawberries? Some people like strawberries, some people don't like strawberries.

Now, you did not come into the world knowing that fact. You now know that fact. And I do not think that most of you learned that fact, like in circle time at preschool. Hello, children. Today we're going to decide if you like strawberries. You like strawberries. You don't like strawberries. They make you break out in a rash. No, right?

You learned about your taste for strawberries in some different way. You learned that you should dive under the bed when you see a flash of lightning because you're scared of the thunder. You learned about that not because you went to thunder school, but because you learned about the association. You picked it up from picked it up from the environment.

And that's what-- it's that sort of association learning that I'm going to talk about today. I'm going to talk about two different versions. One of them is learning associations between stimuli in the world, traditionally known as classical conditioning, or Pavlovian conditioning, in honor of Ivan Pavlov.

The other is learning the associations between what you do and the consequences of what you do, often called operant conditioning, sometimes called Skinnerian conditioning in honor of B.F. Skinner, one of its great practitioners.

Now, let's start with this stimulus form of learning. The basic classical conditioning experiment is one of those experiments that people know about before they come into a psych class, typically, that's Pavlov and his salivating dogs. What Pavlov was doing was studying digestion in-- and dogs were his animal of choice.

In fact, he's also a Nobel Prize winner, but his Nobel Prize lecture-- well, it probably is posted on the web somewhere. But anyway, he won the Nobel Prize for working out the digestive-- work on the digestive enzymes in, I think, the stomach.

He then decided he was interested in digestive enzymes in saliva and he needed a supply of saliva. Where you're going to get saliva from? Well, what he did was he put a little cannula-- a little hole in the cheek of a dog and a little collection tube.

And then, he took his dog, put him in a harness, and sprayed powdered meat into the dog's mouth. If you spray powdered meat into a dog's mouth, the dog does what? The dog salivates. And since the dog's got a little hole in here, some of the saliva drools out and you can go off and study it.

What Pavlov discovered that was of interest to him-- well, it turned out to be of interest to him-- was that he'd get up in the morning, grab the dog, put the dog in the harness, and the dog would start salivating before any meat powder showed up. Pavlov's thinking, I can save on meat powder.

The difference between the rest of us schlumpfs and Nobel Prize winning guys is when they see something interesting, they know that it's interesting. And what Pavlov figured out is the dog is anticipating the meat powder. The dog has learned that this situation means meat powder in some fashion. Gee, that's actually more interesting, said Pavlov, than dog drool. I'm going to study that for a while.

And so what he did was he set up the more familiar form of this where-- let's introduce some terminology. He had food, the meat powder. We will call that an unconditioned stimulus. The characteristic of an unconditioned stimulus is that it produces an unconditioned-- whoops-- an unconditioned response, saliva in this case, without you having to do anything. You get your off the shelf dog, that dog is going to drool for you if you give it food.

Then, what he did was to pair that unconditioned stimulus with a conditioned stimulus, like a bell or a tone or something like that. So now, he was going bell, food, and the dog obligingly drooled. Bell, food, drool. Boom, boom, boom. Over and over again.

Then, the critical thing to do is you turn-- you do just the bell by itself and skip the food. What do you discover? You discover that the bell by itself now produces the saliva. And when the bell-- when the conditioned stimulus alone is producing the response, we call that the conditioned response. The animal has learned an association between bell and food. And you're seeing that learning by measuring this conditioned response.

That kind of learning, which is what Kandel was studying in aplysia, is presumably reflexive, automatic, outside of the realm of consciousness. Even if I fall into sloppy kind of language later, understand that the dog is not presumed to be sitting there thinking, bell. What does that bell mean? That bell symbolizes for me the appearance of-- no. Reflexive. Happening automatically. Doesn't matter what the dog does or does not think.

Indeed, you can fall victim to this thing quite apart from your conscious wishes, desires, whatever. There's a vision researcher a half a generation older than me who's still mad at me because 30 years ago I couldn't escape being conditioned in this fashion.

Back in those days, the apparatus when you were putting up a visual stimulus made a noise. I was supposed to push one button-- as quickly as I could. One button if I saw something and another button if I didn't see something. But every time he put up the stimulus, the apparatus went click. And I just would hit the stupid button.

I had learned the association between click and the stimulus. And my little sea slug brain was refusing to be overridden by this great big conscious apparatus of this guy wants me to look for some really boring thing. And he's still mad at me. It's very, very sad.

OK, this is a very rule governed behavior. In fact, it probably says that on the handout. No, it says, what's being learned here? It's a very rule governed behavior. Pavlov himself worked out many of the rules. And then, a large body of research afterwards continued that effort.

And let me tell you about some of the constraints on this form of learning. And then, explain a bit about how it is that this might actually have anything to do with learning more interesting things than what the sea slug learns is that, gee, when I sense shrimp juice as I recall in the water, that means somebody's going to poke me. So I should retract my gill. That's the response that the sea slug makes is to retract its gilll. Response the dog makes is to drool.

I mean, that's nice, but it's not-- how does that relate to anything like human behavior? So let me tell you a few of the constraints on it. First, a constraint that's really more like an advertisement for next time is to say that we talk about these conditioned stimuli out there, in this case, the bell.

Even something like deciding what the stimulus is is not an entirely trivial business. So suppose I was to come in here with-- well, with a rabbit, let's say. Come here and with a rabbit. Every time I walk in with the rabbit, I activate the electric grid that's in your seat and you all get a little zits and you jump in the air.

Eventually, I bring the rabbit in and what happens? You jump without me having to bother running electric current through your posterior. What's the stimulus there? Well, we naturally would say, it's the rabbit.

And it could be the rabbit's ear. It could be rabbit ear combined with blackboard or something like that. It turns out that we have natural-- and animals have natural ways of carving up the world into stimuli, into objects, perhaps.

And even answering-- when you're trying to figure out what's governing the behavior of an animal, it's not even immediately trivially obvious what that stimulus might be. That's what we'll end up talking about for the next few lectures after this.

But sticking with the realm of the constraints on the learning itself. One important constraint is that you have to notice the relationship between the conditioned stimulus and the unconditioned stimulus. One version of that is completely trivial. If I use a very quiet bell that you can't hear, you don't learn anything about it. You know, big deal.

The more interesting case is illustrated by a phenomenon known as overshadowing. Suppose what I do is I produce-- I have two stimuli that appear at the same time, this bell and a little light. A nice little LED or something like that.

So bell and light followed by food. Bell and light followed by food. Bell and light followed by-- there's a loud bell and a little light. Got it? OK, what happens in this situation is that when you play the bell by itself, you get salivation. When you play the light by itself, even though the light was always followed by the food, you don't get salivation because the light has been overshadowed. That's where the jargon comes from. The light has been overshadowed by the presence of the bell.

What's the critical control experiment in an overshadowing paradigm? What do you need to know here for this to be interesting? Yeah.

STUDENT: You have to test the light.

PROFESSOR: You have to do a training session with some other dog or rat or something with just the light. You have to go light, food, light, food, light, food and show that this relatively wimpy light, if presented in isolation, will produce perfectly fine learning. In the interesting version of the experiment, that's, in fact, the case.

But when you present bell and light together, the stronger bell captures the learning and the weaker light loses. So this is a silly example where I was bringing in a rabbit. If the rabbit was, oh, I don't know, wearing a little silver ring or something like that, you'd learn the association between rabbit and shock but not the association between ring and shock, perhaps, because it was too small to-- it was overshadowed as a stimulus.

It is also critical that the CS predicts the US. Did I get that? Right. Yes, the conditioned stimulus has to predict the unconditioned stimulus. It is not adequate simply for the CS to just show up every time the US shows up.

So as a silly example, everybody who bombs the midterm in this class-- almost everybody. Maybe it's not perfect, but it's a pretty good association. Almost everybody who bombs the midterm in this class drank something shortly before doing that.

So drink, bomb. And it works every year, so there's lots of pairings. Drink, bomb. Drink. Look, I wouldn't on the basis of that give up drinking. There's no predictive value there because you drink a bunch of times when nothing has anything to do with the midterm. There's no relationship.

What's important is a contingent predictive relationship. This is a brain mechanism that's there to learn what in the world predicts what other thing in the world. In the standard Pavlovian set up here, the prediction is perfect. Every time you ring the bell, you get the food. If you cut that down so that the prediction is imperfect, it's correlational, you'll still learn, but you'll learn less and you'll learn more slowly.

So if half the time the bell rang you got food and the bell never rang any other time, the animal would learn that bell more or less means food, but it would be slower and less strong. That's not what I wanted to move.

I've been-- speaking of learning, I've been teaching in this room for years and I've still never managed to reliably use these boards. Always push the wrong button.

All right, so suppose if you've got a good strong bell, food, bell, food, bell, food kind of thing, what you're going to do is, over time, you build up a rate of responding to some asymptotic level. What happens if you change the contingency? Most particularly, what happens if you stop feeding the dog after the bell?

STUDENT: The response will go away.

PROFESSOR: The response will go away, not the dog. The dog doesn't have the option in this particular case. Yeah, the response will go away. That's known in the conditioning literature as extinction. The response will extinguish. The interesting question-- I mean, that makes sense, right? If the contingency no longer applies, why should I continue to respond on the basis of a contingency that's no longer valid?

Did you forget-- you, the dog, did you forget that contingency or unlearn that contingency? Is it gone? The answer is no, it still seems to be there in some sense, it's just you're no longer making the response. How do we know that?

Well, for instance, if you were to give the dog a little break, send him home, then put him back in the apparatus, you would discover that the response-- even if you never presented any more food, the response would come back and then extinguish again. This is known as spontaneous recovery.

And to use the language that we shouldn't be using, it's like the dog saying, all right, bell predicts food. Bell doesn't predict food anymore. I mean, it's not like I forgot that the bell used to predict the food. I'm not a dumb dog. I remember that. But it doesn't anymore, so why should I drool?

OK, now, I wonder what situation we're in here. Are we back here or are we here? Well, we'll drool a little bit just to see. Again, the dog is not presumed to be doing anything like that kind of thinking. But it's not that far off from more complicated situations that you could imagine in the real world.

So you develop a relationship with another person that produces some rate of conditioned response or something like that. You're responding, in any case. Then, she/he decides not to have anything further to do with you. Response disappears. Then comes, oh, I don't know, summer vacation and you see her again after vacation. Do you emit a response?

(IN A HIGH-PITCHED VOICE) Well, are we here or maybe we're back here? Could be-- so maybe a little response.

[LAUGHTER]

That's spontaneous recovery. I mentioned last time that you are better set up to-- the original notion was that you were a general purpose association learner, you could learn anything. The evidence from preparedness, which I talked about before, suggests that you are better prepared to learn some things than others. So you're better prepared to learn that snakes are scary than that bunnies are scary, for instance.

It's fairly general purpose, or at least, there are aspects of association learning that seem to be general purpose, but not all of it is completely general purpose.

One constraint-- one thing that you were certainly set up to do is to notice associations only over a limited time window. So here, oh, I lost it underneath there, right? Tone followed by food. Tone followed by food. That's the CS followed by the US, right? So let's take this asymptotic point here and plot that over here.

Now, suppose that we systematically vary the relationship in time between the conditioned stimulus and the unconditioned stimulus. Well, this is on the order, say, of one or two seconds. If I feed-- if I ring the bell today and feed you tomorrow, how much learning do you figure we get?

Not a lot. First of all, it's a very long, boring experiment since we need lots of pairings. But basically, you don't have to go that far. But it's going to fade off in this direction. You can't have too long a period between the conditioned stimulus and the unconditioned stimulus.

0 doesn't produce much conditioning and negative relationships. So I give you food. And then, I ring a bell that says food. That was food. Food. That was food. OK, little-- it doesn't hit 0 at 0, but there's really basically very little happening down that way. So there's a narrow time window within which this chunk of you is looking for associations out there in the world.

The real world example here might be bump signs on the highway. When is a bump sign on the highway useful? If the bump sign shows up at exactly the same place as the bump, thank you. You want it to show up a little bit before the bump. It shows up after the bump. That was a bump.

[LAUGHTER]

It's not that useful. So it's limited constrained in time. Now, how can we go and get from this nice rule-governed behavior to more complicated behaviors? One route to that is-- or one way to see how that might work is a paradigm in the conditioning literature known as sensory-- did it again, didn't I? Look at that. Got it wrong on both dimensions. It's called sensory preconditioning.

I'll just restart this over here. It's easier. Here's how a sensory preconditioning experiment might work if you were a rat or a dog or something like that. So step one, I'm going to show you a light. And then, I'm going to ring a tone. And I'll do that a bunch of times. Light, tone, light, tone, light, tone. I'm going to play the light alone or I play the tone alone. Do you salivate? No, why should you salivate. It's got nothing to do with salivation.

Step two, tone, food. That's just the classical conditioning paradigm. Tone, food, tone, food, tone, food. Play the tone alone. Well, tone, food, whoops. Saliva. So now, I'm going to play the tone alone. And lo and behold, we know I'll get the conditioned response of salivation here.

The critical test is now I play-- I show the light in isolation. What happens? Well, what happens is that you get the conditioned response. The animal will now salivate in response to the light alone.

Why is this interesting? Well, what's the animal doing? Well, again, we can think in conscious terms that aren't appropriate. The animal is saying light. I know about light. Light means tone. Aha, tone. I know about tone. Tone means food. So if tone means food, I should be salivating. And if light means tone, I should be salivating there, so I think I'll just start drooling all over the place now.

What's important about this is you can see-- actually, you can see graphically here, I'm managing to get more remote from the response. I'm working my way backwards. Suppose it was the case that at the end of every one of my lectures around 3:23, I was to say, in summary. And that your habit was to immediately run out of class thereafter and get yourself a snack.

You might discover that as I said, in summary, that somebody walked in. You might discover that as I said, in summary, that you started to feel hungry because of a learned association. It might never occur to you until you examine this why you feel hungry when I said, in summary.

But if we look at this timing thing, it's not the case that I say, in summary, and one second later I blow M&Ms into your mouth or something like that. You've got to somehow explain how you work backwards from the eventual M&Ms here back to something that I said. And this is one way to do it.

You imagine that what you've got is an automatic association learner that's-- its job is to look for if a then situations in the world. And that it is capable of chaining those together. That you can go if A then B, if B then C, if C then-- oh, hey look, if this, then this thing way down the line somewhere.

And if you can start chaining these simple associations together, all of a sudden, you've got the possibility of having a much more complicated kind of-- a tool that could potentially explain much more complicated kinds of behavior.

Now, this is-- so that's classical conditioning. Good for explaining or good for talking about relationships between stimuli out in the world. The other thing-- another thing that you would want to learn about is the relationship between what you do and its consequences.

That's what we talked about a bit last time when I did that puzzle box-- the Thorndike puzzle box example. Thorndike's cat wasn't learning about-- was not so interested in learning about relations of stimuli in the world. What it was interested in learning is if I do this, then I get that fish. And that's the core of operant conditioning.

Now, if this was-- I took this course once upon a time. But I took it at Princeton, where I was an undergrad. That allows me to encapsulate the basic difference between a Princeton undergraduate education-- oh, look at that. There's a guy wearing a Princeton sweatshirt right now.

[LAUGHTER]

Don't cover it up. You just-- unless she's no longer talking to you. That's the extinction curve over there.

[LAUGHTER]

Anyway, where were we? Oh, yes, the difference between a Princeton education and MIT education is that, at Princeton, this course satisfied the science requirement.

[LAUGHTER]

You laugh. But if I was to give this lecture at Princeton and say, at MIT, this course satisfies humanities literature kind of requirement, they'd all laugh. So the main difference between the courses was not in what the guy standing up front did in the way of lecturing.

But here, of course, to satisfy to make it a HASS course, you do a lot of writing. And at Princeton, we had-- it was a lab course. And it's a great pity-- in some ways, it's excellent that you are off writing. But it's a great pity not to have a lab component to this because this operant conditioning thing is a lot cooler if you got a pigeon of your own, which is what we had.

So if people-- you don't end up with a salivating dog of your own. But pigeons, pigeons are OK. So the classic-- the version of an operant conditioning experiment that people tend to know something about before they come into a course is a so-called Skinner box, which is a pigeon or-- well,

[LAUGHTER]

He can be a rat too if you want.

[LAUGHTER]

It's just a matter of where you put the whiskers. Anyway, so the idea of a Skinner box was you'd have an environment that could really control the options available to the pigeon or the rat. And what you'd have in a pigeon Skinner box is a box that had nothing much in it except for a key here that is hooked up to a little micro switch or something that the animal could peck at and we could record the pecks and a little bin down here where the bird seed could show up.

And you stick a pigeon in there, and this is a hungry pigeon because you haven't fed it. And so yeah, it's a bummed out pigeon. But if he pecks there in the right ways, we'll feed him.

So this is simply the somewhat higher tech version of Thorndike's puzzle box you'll recognize. The pigeon has to figure out to peck at the key. If you just stick a pigeon in a Skinner box, even a hungry pigeon, the pigeon just sits there. I mean, you know, does stuff, but the pigeon doesn't immediately say, hey, I take an intro psych. I know about this. I'm going to go and peck that thing.

You have to do what's known as shaping its behavior. The first thing you have to do that's important is you have to get your pigeon from the basement to the fourth floor where the lab was. And should you ever need to do that, the important thing to know is that plastic juice containers are really good. Take the pigeon, you stick him head first in the plastic juice container, not with the juice in it.

[LAUGHTER]

Then, the pigeon is actually quite calm under those circumstances. If you decide how tough can a pigeon be and you don't need to do that because you can't find your juice container, anyway, you just grab your pigeon because you're bigger than the pigeon, try to carry it upstairs, what you end up is with loose pigeons on the second and third floor. So things that you're going to miss out on just because you're writing papers instead.

What you do, once you stick that pigeon in the Skinner box, if you just sit there, it's going to be a long day. What you have to do is shape the pigeons behavior. It probably says shaping somewhere on the handout.

Yeah, use the law of effect to shape the animal it says there. That doesn't mean like--

[GURGLING SOUND]

What it means is that the pigeons just sitting there moving around doing pigeony like things feeling hungry. And the pigeon turns towards the wall that's got the key on. And you push a little button that gives them a little birdseed. And it goes--

[GURGLING]

Eats the birdseed. And goes back to doing pigeony things. But now, the law of effect is working. That birdseed is a positive reinforcer, right? So the chunk of the bird's brain that is doing this association learning between action and its consequences saying, birdseed. That was good birdseed. I liked the birdseed. What do I do to get more birdseed?

Well, it's not saying that explicitly. But it's saying, we'll do whatever we were just doing, OK? Well, now you don't want him to just be looking at the wall. So now, after a couple of rounds of that, you say, OK, bird, you only get the bird seed if you move a little closer to the wall.

The bird moves closer to the wall. And you say, all right, now you only get it if you're looking at the key. Now you're only getting it if you're right up there. And eventually, you get the bird pecking at the key. And then, you can go off and do other cool experiments from there. But you have to shape the behavior in the way you want. And the key business, I mean, there's a key, there's the bird food, that's somewhat arbitrary. It's just what Skinner wanted to use to control the bird's-- to make it possible to measure the bird's behavior.

We had a very clever pigeon when I did this lab. And we got through the stuff that was in the lab book really quickly probably because the pigeon had been the pigeon for a half a dozen other-- oh yeah, I'm supposed to fake it here and claim I don't know what's going on here. All right, oh good. We're shaped now. Give me the food.

Anyway, we decided that this-- we still had some time. We thought, we'll mess with this pigeon. And so we started reinforcing the pigeon for turns. So no bird seed unless you make a quarter turn, OK? No quarter unless you make a half turn. We eventually had the bird making three full turns for each-- and then staggering over to the bird seed. But we had this great ballet dancing pigeon.

And that's, in fact, the basis of trained animal acts that you may have seen at various places. And I should advertise that it's not confined to pigeons and things like that. You can perfectly well condition human beings. I think I used in the last lecture the example of jokes as a case of behavior getting shaped. Jokes that get rewarded with laughs or get repeated. Jokes that get rewarded with smacks don't get repeated.

But again, going back to my intro psych class, I should tell you that it is possible to condition the professor. What you need to do is figure out, well, what reinforces professors? The answer, at least in this sort of a setting is, students who look kind of interested and smile and write notes. At least, that's what my professor who was, in fact, an important learning theorist in his own right, that's what he told us.

He said, if you look-- it's reinforcing if you're looking like you're interested and you're writing notes. And so you should, he told us, be able to pick a behavior that the professor's doing and reinforce it by smiling and writing notes at the appropriate time.

So I was going to be a psych major. I did what I was told. I was sitting about there. My friend and I, every time-- the next psych lecture, every time the professor came in, moved to the left, we smiled and took notes.

[LAUGHTER]

And by the end of lecture, he was most of the way out the door.

[LAUGHTER]

It was great. Now, as evidence for the notion that this is quite unconscious and reflexive kind of learning, we told him about it afterwards. And he absolutely believed that it was possible that it had happened and absolutely reported no awareness that it had happened.

So you're welcome to try this. You're even welcome to try it here. Let me tell you that I have made this offer in years past and nobody has reported a great success. But the great failure was when one recitation section got together and decided there's strength in numbers.

But it turns out that if everybody in a recitation sits, I still remember, they were all sitting there. If y'all sit in one place. And I don't remember what I was supposed to do, but when I did it they all went--

[LAUGHTER]

I kind of cracked up, but I don't believe I learned anything in particular. But it can sit out there as a challenge for you. If anybody gets it to work on the math guy or something. Is Arthur Maddox still teaching 18 0 whatever?

[INTERPOSING VOICES]

PROFESSOR: 1803 in the spring? Oh, that's too long to wait. I've never been sure Maddock would notice anyway, but that's a separate issue.

[LAUGHTER]

Anyway, if anybody succeeds in conditioning their professor, do let me know.

The point here is that this is a form of learning that does apply in the human world as well as in the pigeon world. Once you've got that animal shaped to do what you want, you can then start studying the rules that govern that behavior.

Let me say a word about schedules of reinforcement since I see that it's there on the hand-- actually, it doesn't say schedules. Oh it does. Fixed and variable ratio and fixed and variable interval schedules of reinforcement.

What that means is if you're reinforcing the pigeon every time he pecks, he's on a fixed ratio one schedule. You shape up the behavior. And at some point, the pigeon is trained. And if you're reinforcing him every time he makes a peck, he will emit pecks-- so let's call this the peck rate.

He'll emit pecks at some asymptotic rate, not unlike the salivation thing over there. Actually, I want to save myself a little space here so let's only go this far. So he has an asymptotic rate of responding.

Now, suppose that you change the rule. Suppose you say, you don't tell-- well, you can say it to the pigeon because the pigeon won't understand you. Say to the pigeon, OK, now you get reinforced only every 10th peck. What's the pigeon going to do do you think? How is the behavior going to change?

It's going to go squiggly apparently. What is this? No, every 10th peck, not every ten seconds. I'm interpreting your squiggle as the answer to a later question. So save the squiggle. It'll be good.

STUDENT: Peck faster.

PROFESSOR: Peck faster. Whoever said that. Well, if you've got to work harder to get the same reward, you're going to end up pecking faster, right? So you think about it in terms of your-- once upon a time, you could get an A in whatever, writing for writing like what my eight-year-old wrote this morning. Seven sentences of at least six words each. Took him all of breakfast to do that because he'd lost his book. And anyway, it's a long story.

So he emitted his writing behavior to a relatively low rate. If you emit your writing behavior at that rate, you guys are going to be in serious trouble, right? To get the same reward, you now have to produce more work. And as a result, you'll crank up the level.

So that would be a fixed ratio-- what I was describing is every 10 pecks that would be a fixed ratio of 10. A variable ratio is, on average, every 10 pecks you get reinforced. And that, actually, produces even higher levels of responding.

Now, what the squiggle was the answer to is what happens if instead of a ratio schedule you have an interval schedule. Suppose you reinforce the behavior only for the first peck after every minute. Make a different graph over here.

So bird can peck it all they want, but only when they peck after this interval marker goes by are they going to get reinforced. Well, I mean, a pigeon with a wristwatch is going to look at his watch, and say-- pigeon doesn't have a wristwatch, of course. But they do have an ability to estimate time.

And so what they do immediately after getting rewarded, they just sit there. So this is, again, going to be rate of response. They just sit there. And then, they say, is a minute gone yet? Oh, a little. No, nothing happened. Got to be now. No, it's still not. It's got-- and it's got to be now.

[BANG]

[LAUGHTER]

And you get this very scallopy behavior. And now, if you don't think that has anything to do with human behavior, well, first of all, you can ask yourself about your writing output. So for instance, there's a paper due on a regular interval.

[LAUGHTER]

A smart pigeon would be emitting words at about that rate, right? Well, actually, I did meet a student who claimed he was doing that. More power to him. But I think most of you this-- when it's due Friday, I think this point here is Friday 4:00 PM or something.

Anyway, so people do-- or you come to my house where, if you're an eight-year-old, you get your allowance on the weekend, right? But only if you do your chores. So Monday is my kid asking what chores he should do? No. Saturday morning? You bet. What can I do? What can I do? I made my bed. Can I lick the floor now?

[LAUGHTER]

So these sorts of things-- oh, and the downside of this is-- one of the reasons for paying people on a Friday is that if you pay them on a Thursday they don't show up on Friday. Even though, at some level, they're much more likely to be sick on the Friday if you pay them on the Thursday. So you want to pay them on a Friday and then they can be sick all weekend on their own time.

So that's what we mean by schedules of reinforcement. This is also-- operant conditioning is also a very rule-governed behavior. It is governed by rules that are in many ways similar to the rules for classical conditioning. So unsurprisingly, let's just do this one as an example. It's probably later on in the handout.

So I reward little old pigeon every time he pecks the key. Now, I stop rewarding him. What happens to the behavior? [IMITATES BOOM]. I take that [IMITATES BOOM] to mean, yeah, it extinguishes. If I give him a break, I'll get some spontaneous recovery and it'll extinguish again, just like you'll get in classical conditioning.

What happens in these? How does the extinction compare in a schedule like this to a schedule like this? How about a hand? I'm hearing good mutterings, but nobody has a hand. Yeah, I know-- well, all right, so slower. I heard some more slower mutterings.

It's slower. If that's not intuitively obvious, ask yourself about the difference between a slot machine and a Coke machine, right? You put a buck in the Coke machine, you get a Coke, right? You put a buck in the Coke machine, you don't get a Coke. How many more bucks do you put in?

[LAUGHTER]

Well, all right, so you're not too dumb. That's good. OK, now you go to Las Vegas. You put a buck in the slot machine and you don't get anything out. What do you do? Another buck in. Actually, Las Vegas is a beautiful example of applied conditioning theory. We could send everybody on a class trip that, actually, we did that, didn't we?

There's that bestseller about-- that was more applied math than applied psychology. The applied psychology thing is when a slot machine pays off, traditionally, what happens is the coins come out, right? Not now. Often, they're electronic, but they simulate it anyway. The coins come out.

What do the coins hit when they come out?

STUDENT: Metal.

PROFESSOR: Metal. Why do they hit metal?

STUDENT: Make a sound.

PROFESSOR: To make a racket. Why do you want to make a racket? Well, that racket itself becomes positively reinforcing. You want to hear that sound, right? You want to put that buck in and hear--

[CHING]

Oh yeah, this is very good stuff. Now, slot machines, when you-- I mean, you don't go off and play slot machines a lot but you see the movies. I hope. Anyway, slot machines are typically located in small soundproof booth. No, where are they located?

[INTERPOSING VOICES]

PROFESSOR: On a giant floor with thousands of slot machines. What does that mean? That means when somebody five blocks over gets a payoff, what do you hear?

You hear the sound. Well, that sound is what you're conditioned to-- you've been conditioned to find that reinforcing and you are so dumb.

[LAUGHTER]

They know that you're willing to be in reinforced by somebody else getting paid off. And you say, oh man, there's money happening here. Let me put some more into this thing. Anyway, lovely applied psychology in the worst sense of the word. Well, maybe not the worst. I can think of a few others.

But while I'm thinking of a few other worse ones, take a minute to stretch. Reinforce your neighbor. Wake up your neighbor if that's necessary.

[SIDE CONVERSATIONS]

STUDENT: You're not stretching.

STUDENT: Oh, sorry.

PROFESSOR: What are you reading? Anything good? Oh. This is good, I suppose. Is it suitably high? Is the price suitably high?

[SIDE CONVERSATIONS]

PROFESSOR: Great. So you know sleep deprived and depressed all at the same time.

OK, let me talk. There is a sense in which you are a general purpose association learner and have mechanisms in your brain that are basically saying, teach me about the associations between stimuli or the associations between my acts and their consequences.

And I can be fairly broad minded about what those stimuli are or what those acts might be. And then, there are special purpose versions of these things designed for very particular tasks.

Let me fish around for what's perhaps the best example of this. Is there anybody here who knows that-- who can think in their life a food that they used to love but they now hate or won't touch and they know exactly why. Yeah, out there in the cheap seats.

STUDENT: I ate a chocolate frosting cupcake when I was little and I went home and I got the flu and I don't eat it anymore.

PROFESSOR: Perfect. Gee, first try. No weird examples this year. Now, you got the flu. Let's get a little graphic here. You got the flu and what happened?

STUDENT: I vomited a lot.

PROFESSOR: Thank you. That's what we needed to know. And not only that, we also need to know, do you believe that the chocolate cupcake gave you the flu?

STUDENT: Not any more, but I used to definitely think that.

[LAUGHTER]

PROFESSOR: Yeah.

STUDENT: All right, when I was in--

PROFESSOR: Oh, he doesn't care that I already have a good example. He wants to give me another cool example. All right, all right.

STUDENT: It's just another one.

PROFESSOR: All right, it's not as good?

STUDENT: When I was in kindergarten, first day of school or the first week I ate lunch in the cafeteria at school. And I got so dreadfully sick. And I didn't eat a school lunch again for five, six years.

PROFESSOR: This is why he's slim to this day.

[LAUGHTER]

Yeah. Actually, we'll come back to that example. It serves a useful function too. Now, the important pieces of this-- so she ate the cupcake, she got sick, she knows at some level it wasn't the cupcake that made her sick. But there is a chunk of brain that says, I got sick. I ate cupcake, shortly thereafter, got sick. Don't want to eat cupcake anymore.

This is a chunk of your brain that is quite immune, or it's actually very hard for you to talk to and say, wait a second, we used to love cupcakes. Cupcakes are really good. It doesn't want to hear about this. It says, you ate it, you got sick.

And this is a form of association learning. It's got some-- well, the original version of this was discovered in-- well, the original version was discovered by some rat who ate a cupcake and threw up. But in the learning literature, it was studied by a guy named Garcia. So it sometimes goes by the name of the Garcia effect.

What he did was he had rats in a Skinner box situation. And he gave them a new flavor of water like spearmint or something. And then, he made the poor rat sick. Strapped it to the turntable of a record player.

[LAUGHTER]

An experiment you can't do anymore because you can't fit the rat into the CD drive.

[LAUGHTER]

And then, the next day, you give the rat a choice. You want spearmint water or you want this other new weird flavor? And the rat goes for the other flavor. Says, forget. I'm not going for the spearmint water.

But this is a very specific form of learning. It's sometimes called taste aversion, which is not quite right if you happen to be a sensation and perception person. Taste, as the sensation that your tongue does for you, is really restricted to sweet, sour, salty, and bitter.

What you become averse to is the flavor, most of which is smell. Or perhaps, in human cases, a larger context like school lunch or something like that. But there are strong restrictions on what it is that you will develop an aversion to, particularly if you're a rat.

So the other version of the experiment that Garcia did was he gave rats a new bottle of water, and this one was flashing lights, right? Water never did that before. Drink the water, get sick. Next day you got a choice between the flashing light one and the one that's making clicking noises or something like that. Rat doesn't care.

This little chunk of brain was not built to figure out that flashing water is a problem. This is a chunk of brain that is there to tell you that there are things-- it knows some things about food. Foods have flavors. And if you get sick so chunk of time after eating, you should avoid that flavor the next time and things associated with that flavor.

Also, the timing-- where did my timing parameter go? The timing is different. If you eat something and get sick instantly, nothing happens. You don't learn anything. Because this little box knows-- it takes a while to get sick. I think your story was, ate the cupcake, went home, got the flu. Hours later she's losing it and this little chunk of brain is working back in time. It's got its own time window.

It's working back in time saying, what did we eat a few hours ago? Cupcake. Don't want to eat the cupcake. Very useful thing, right? If you're a grazing-- an omnivore kind of animal, very useful thing to have a device that says, don't eat it because last time you ate it, it made you sick.

Oh, the other thing, when we were doing this, tone, food, tone, food, pairing, over and over again. How many cupcakes did you eat?

STUDENT: One.

PROFESSOR: One cupcake. And you repeated this lots of times? No, right? Because this is a little association mechanism designed to keep you alive. And if you eat the cupcake, almost die, you don't want a bunch more pairings of eat the cupcake, almost die, eat the cupcake, oh man, that time, dead.

[LAUGHTER]

So one trial learning, which by the way, when Garcia discovered it was thought to be impossible. But it's a special case of an association learning that works in a single trial.

It has its pitfalls. And in fact, it's cognitive impenetrability. The fact that it doesn't care what you know is a problem. So except for the seniors here, of course, nobody here has ever-- alcohol has not passed your lips. And even the seniors have not drunk to excess.

So what may be news to you is that if you drink too much alcohol, you can get sick.

[LAUGHTER]

They knew that. OK, anyway, drink too much alcohol, you get sick, right? The question is, what do you develop the aversion to? A smart system rather like this-- the writing example. The smart system would say I just drank 12 screwdrivers in a row or something like that. No more vodka for me.

The problem is that alcohol has very minimal flavor to it. And what happens is you drink-- orange juice plus vodka or something like that, you get violently sick. I'm not going to touch that orange juice anymore.

[LAUGHTER]

A pile of beers and a pile of pretzels, you're violently ill and the pretzels just look miserable to you. Yeah, so it is a disadvantage. It's clever that this thing is out there to save your life. Unfortunately, in civilization, we have figured out clever ways to misuse it, I'm afraid.

Well, since it says example two is superstitious behavior, let me say a quick word about superstitious behavior. This is the best time to see what Skinner called superstitious behavior because all you have to do is watch baseball.

Law of effect says, if something good happens, you're going to do what you were doing just beforehand. Think about baseball. What's something good that can happen? You get up to the plate, you hit the ball, it goes over the fence, and you get to run around and stuff.

Well, what were you doing just before? Well, maybe you wiped off your shoes or something like that. So next time you come up to the plate, you wipe off your shoes, don't think anything about it. Pretty soon, you're doing this all the time.

Some sportswriter asks, how come you wipe off the shoes every time? The plate is slippery or some story like that. You've got no idea why it is. But it's probably been shaped into place by the contingencies of reinforcement.

The second best place to see this, in my experience, is at exams. Look around at exams, perhaps, look at yourself at exams and say, look at those three neatly lined up pencils. Points exactly lined up there. The Coke has to be exactly to the left of the test paper and the Snickers bar has to be right there or I'm going to flunk.

[LAUGHTER]

Again, if you're asked about it, the answer is, I don't know, it just makes me feel comfortable or something like that. But it's, again, superstitious behavior perhaps shaped into place by this law of effect kind of work.

The schedules of reinforcement. Let me say a bit more about how this applies out in the real world. Schedules of reinforcement are the sorts of things that we do to ourselves all the time. For example, in child rearing, is that actually the next example on the handout maybe? Yeah, it says something about parents and children.

So think about-- well, suppose that you're a little kid, which you once upon a time you were, and you want a cookie. So you say, can I have a cookie? And your mom gives you a cookie and that's good. The next time you say, can I have a cookie? And it's just before dinner, so she says no. So what do you do?

Well, maybe you extinguish the behavior immediately and you're a perfect model child and you end up at MIT or something like that. But odds are that you-- can I have a cookie? Please can I have a cookie? Pretty please can I have a cookie?

Eventually, you're mom, oh here, have a cookie, right? All right, so what have you done? We've now moved you from an FR1 to an FR something else-- or actually, it's probably a VR schedule of some sort. So now, you're going to be emitting cookie behavior at a much more rapid rate, right?

Can I have a cookie? Can I have a cookie? Please can I have a cookie? I want a cookie. I'd like a cookie right now. Eventually.

[LAUGHTER]

Eventually, your parents get tired of this they decide they're going to cut you off here. But you're now up here somewhere. How does that extinction curve look? Oh man, it goes out to about age 23 or something like that.

[LAUGHTER]

Plus, parents are not good at this. So what you do is, can I have a cookie? Can I have a cookie? Can I have a cookie? Can I have a cookie? I really want a cookie. Oh, can I have a cookie? OK, now we're on a VR 2064 kind of schedule and everybody's going nuts.

The other place you see a similar example is sleeping through the night. The kid cries, right? You go and comfort the kid, put the kid back down to go to sleep. Eventually you decide, this little monster needs to sleep through the night because I need to sleep through the night.

So I'm going to not get up when he cries. He's crying. I can't sleep. It's been crying for like an hour.

[LAUGHTER]

Oh, maybe just this time. All right, now you've rewarded him for crying for an hour rather than for crying for a couple of seconds. And again, you end up with a similar problem.

My eldest actually had this to an absolute art when he was little. Of course, he didn't know it, again, because all nice unconscious and reflexive stuff. He came into the world built so that if he cried for more than a few minutes, he also threw up, right?

It's time for him to sleep through the night. He's crying and you're saying, oh man, if he doesn't go back to sleep in a second, not only am I going to go nuts, but I'm going to have to do the laundry again.

[LAUGHTER]

But if I reward him for-- and you know, so he's-- I don't think they sleep through the night yet. Is this encouraging yet? Mara is getting so depressed there.

All right, you're not babies anymore. Let's give an example of the relevance of this that might have something to do with something closer to adult behavior. One of the interesting realms where these schedules of reinforcement may have an unintention-- well, none of it has an intention to it-- an unfortunate negative consequence is a phenomenon known in-- at least in the research on sexual behavior as getting to yes. What is this about?

Typically, in heterosexual relationships, the request, verbal or otherwise, for some positive reinforcement of a sexual nature comes from the guy and is delivered to the woman in some fashion. Typically, that request, verbal or otherwise, is met with a no initially.

If the relationship develops over time, of course, eventually, odds are that there will be a yes. Well, what does that sound like? That's some sort of a variable schedule of reinforcement?

Oh, why might that be the case? We'll talk about this more later in the term. But there is interesting evidence that guys are not just some unbridled bag of hormones that-- but that they actually fall in love faster than women do.

Well, it's the Romeo and Juliet thing, right? He's allegedly in love with Rosalind. He goes and sees Juliet. And two seconds later it's, oh, she doth teach the torches to burn bright and all that stuff. You know, woo, he's gone.

[LAUGHTER]

And when he walks out of the party later and meets his friend the psychologist who does a survey kind of thing, he will report he feels like he's in love. And I don't know what-- Juliet may be a different example. I mean, they all end up dead and things like that.

[LAUGHTER]

But typically, she will be slower to report that she feels like she's in love. So he's sitting there saying, I'm in love and I'm going to be with her forever. And if I'm going to be with her forever, there's certain things that we might be interested in. And so maybe-- and she's saying I just met him five minutes ago. I mean, what is this about?

[LAUGHTER]

Anyway, the result is a some sort of a variable ratio schedule of reinforcement. And in most cases, this works itself out just fine. But suppose the relationship doesn't go particularly well. Suppose, for example, well, where's my extinction curve? Here's a good extinction curve.

So you're up here somewhere on this variable ratio schedule of reinforcement. And let's say she decides that this is not a relationship that she's interested in anymore. It takes him now an extremely long time, potentially, to get the hint and stop pressing that key in the Skinner box.

Well, we'll talk later about the ambiguous nature of communications in this regard. If he's not sure whether she's saying yes or no or whatever, you can end up in trouble. And this is an example of applied where these rules of association learning suddenly become relevant in behavior that's much more complicated than whether or not your pecking a response key in a Skinner box.

It's that kind of application, the possibility of going from these very simple behaviors to much more complicated, much richer behaviors, that drove what was the leading school of psychology in the first half of the 20th century in the US, which was known as behaviorism, which was the doctrine that basically said, look, classical conditioning, operant conditioning, a couple of other bits, those are the atoms of behavior.

And we can build up the molecules and the rest of the organism out of those atoms. That's all we need. In the same way that-- and so I put this quote from John Watson, who's the founder of American behaviorism, on there where he says, we can write a psychology-- define it as the science of behavior. That was a move by itself, right? No science of behavior and mental life or no science of human mental life or something.

The science of behavior, of observable behavior. Never go back on our definition. Never use terms like consciousness, mental states, mind, content, and so on. It can be done in terms of habit formation, habit integration, which are terms for this association learning.

He is basically saying that, look, if you're taking chemistry and you go to the chemistry prof and say, these elements you've got here, they're nice elements, but I think I want another 16 kind of constructs to explain what's going on here.

The chemistry professor is going to look at you and say, no, sorry. These are the elements. Everything can be built out of this stuff. You don't like it, you got a problem.

And the behaviorists were saying essentially the same thing about psychology. They thought that the very small set of atoms, in this case, these laws of association were the elementary properties. And you could build up everything out of that.

That carried with it a couple of interesting bits of ideology that will recur later in the course. One of them was the idea that you were, basically, an association engine. You were there to learn associations. And that everything that was worth knowing in psychology, everything that was worth studying in psychology was something that was learnable. That it was not interesting, beyond-- for their purposes, kind of trivial things like you came with eyes and not radar dishes.

Beyond that sort of thing, It was uninteresting what nature had provided to you. Psychology was what-- was the effect of what the environment-- the interaction of the environment and these atoms of learning that everybody was essentially identical until this learning started.

And what got you to MIT was the contingencies of reinforcement over the prior 18 years of your life. This resonated with a certain Democratic current in American political thought. In grade school, civics lessons form. This was the doctrine that anybody could grow up to be president. You didn't actually have to be a Bush to be president. Anybody could be president.

What was important was how the environment shaped you. That's a doctrine and it's radical in its strongest form known as empiricism. That look about right? Empiricism.

OK, the opposite pole is nativism. The notion that your genetic innate endowment is determinative. And that you're here because you were born with-- you've got the gene for calculus in there somewhere. Some of you may have now decided you lost it somewhere. You had it once upon a time.

So the notion here is that it's-- what's critically important-- what the behaviorist thought was they were strongly weighted towards this empiricist environmentally driven side of things.

In a couple of minutes, why aren't we-- why am I using the past tense? Why aren't we all behaviorists now? There are a number of reasons. I suppose the broadest reason is, look, a course like this is full of on the one hand position and on the other hand, empiricism and nativism.

And I'll give you a hint for the exam. The answer always turns out to be both, right? The extreme positions, the extreme theoretical positions in psychology never seem to quite work out.

Less trivially, look, the association-- the basics of association learning theory are absolutely worth a lecture in this course and absolutely worth chapter 4. And you'll find on the handout some guide for reading in chapter 4.

But they're not all of psychology. Watson's statement says, we can build a science of behavior without ever talking about consciousness, for example. But that's an interesting topic.

And if I want to know about it, I don't want to be told by the field's ideologues that it's not a legitimate area of inquiry. I don't want to be told that emotion, that the feelings are-- that the feeling aspect of emotion is not a legitimate topic of psychology and that the only thing we can do is observe how many tears are coming out or something like that.

It proved to be an incomplete psychology. In the same way, I should say, that the occupants of the nativist poll at the moment are the most dogmatic of the evolutionary psychologists. Evolution has wonderful explanatory power in psychology. But evolution is not itself a psychology. It's not a complete psychology by itself. It doesn't give you the richness of the field as a whole.

Another important reason we're not all behaviorists is it turned out to be somewhere between uninteresting and wrong to argue that everything that was of any interest was learnable. Language, we'll see later in the course, is an example. Of course, you learned the language that you speak. But you learned it because you had a brain that was built to learn a language.

And that innate endowment, that language learning endowment, is worth study in its own right and that wasn't something that fit into the behaviorists' program. So that's why you're not a behaviorist now. But even if you're not a behaviorist, you still want to read chapter 4.