LING 1010

Language and Mind

Prof. Jon Sprouse

04.26.21: Linguistics, Epistemology, and the Philosophy of Science

Epistemology

You didn't even know it, but the first two sections of this course were a detailed case study in an area of philosophy called **epistemology.**

Epistemology is a word of Greek origin meaning "study or theory of knowledge". The field of epistemology is concerned with two broad questions:

1. What is knowledge? What does it mean to "know" something?

2. How is knowledge acquired? What is the process by which humans come to "know" something?

You might notice that these are the two driving questions behind the first two units of this course!

What is knowledge of language?

What we've learned is that to know a language is to know the mental representations of your language.

And we have seen that knowing the mental representations of your language means that you know the the units (phonemes, morphemes, syntactic categories) and grammatical rules (phonological rules, morphological rules, phrase structure rules, transformations) of your language.

In short, knowledge of language is the knowledge of the grammar of your language.

And that grammar is complicated:

Units	**Grammatical Rules**
phonemes	phonological rules
morphemes	morphological structure building rules
syntactic categories	phrase structure rules and transformations

How is knowledge of language acquired?

In epistemology, knowledge is divided into two types: a priori knowledge, and a posteriori knowledge:

a priori knowledge: knowledge that comes before experience

Philosophers tend to focus on something called “analytic truths”, which are truths that are true by definition of the words in them. “A bachelor is a man” is true by the definition of “bachelor” and “man”.

But in linguistics, we’ve seen that there is another type of a priori knowledge - the knowledge that is specified by your genes, which may help facilitate language learning.

a posteriori knowledge: knowledge that comes from experience

Philosophers tend to focus on something called “synthetic truths”, which are truths that are established by observation (or in other words, truths that could have been different), such as “ravens are black”.

In linguistics, we’ve seen that evidence from experience must be part of language learning.

From Linguistics to the Philosophy of Science

Infinite sets and Scientific Theories

Recall from the logical problem of language acquisition that we can define an infinite collection of objects - an infinite set. This is easiest to see with numbers.

The set of all numbers greater than 2

Recall from our linguistic theory that the set of sentences in a given language is an infinite set.

The set of sentences in English.

It also turns out that many scientific theories can also be formulated as infinite sets.

All ravens are black.

If your theory says that all ravens are black, that means if you encounter an infinite number of ravens, they will all be black.

Learning infinite sets

Recall that we saw that there is a major logical problem with learning infinite sets from a finite subset.

The set of all numbers greater than 2

2, 4, 8

That logical problem was also an issue for learning a language.

The set of sentences in English.

The cat...

The same logical problem holds for establishing the truth of scientific theories from evidence.

All ravens are black.

The theories are infinite sets, the evidence is (always) a finite subset, so the same logical problem arises in science. **How do we guarantee the truth of our (infinite) scientific theories from finite evidence???**

Finding a process for establishing the truth of scientific theories

Why should we care about the process?

Human advancement has not been steady. Our knowledge grew slowly for thousands of years, and has exploded after the last couple of hundred. The reason is **the scientific method**.

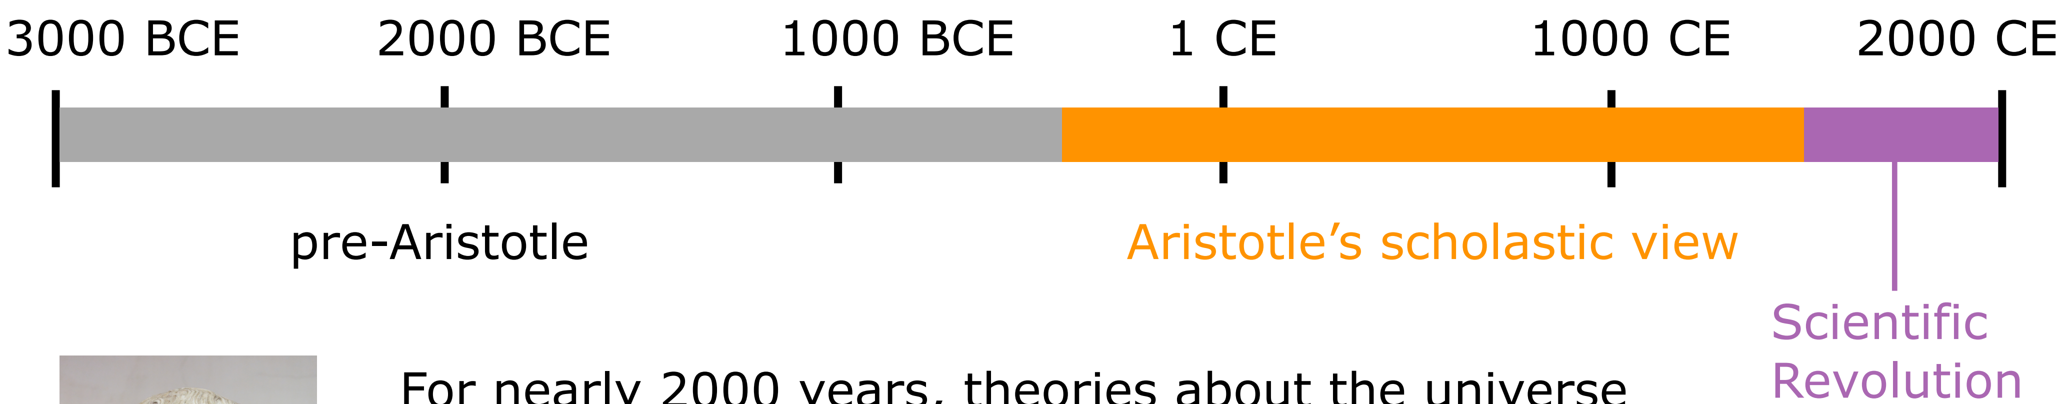

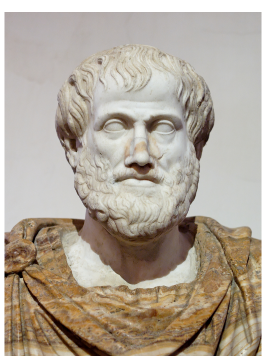

For nearly 2000 years, theories about the universe were dominated by Aristotle’s view that there were 4 elements (earth, air, water, fire), and the earth was the center of the universe.

Then from 1550-1700 prominent thinkers started to question those theories, and even question how to build a theory. We call this the age of enlightenment or the scientific revolution. From that point forward, the expansion of human knowledge has been dramatic!

Why should we care about the process?

Scientific debates are becoming more and more relevant to the world we live in. As information becomes easier and easier to access, it is critical that we understand how to use evidence to prove/disprove theories.

The Impact of Vaccines in the United States

Disease	Baseline 20th Century Pre-Vaccine Annual Cases	2008 Cases*	Percent Decrease
Measles	503,282	55	99.9%
Diphtheria	175,885	0	100%
Mumps	152,209	454	95.7%
Pertussis	147,271	10,735	92.7%
Smallpox	48,164	0	100%
Rubella	47,745	11	99.9%
Haemophilus influenzae type b, invasive	20,000	30	99.9%
Polio	16,316	0	100%
Tetanus	1,314	19	98.6%

*Provisional. Widespread use of vaccines in the United States has eliminated or almost eliminated infectious diseases that were once terrifying household names. Credit: Morbidity and Mortality Weekly Report, Centers for Disease Control and Prevention, 4/2/99, 12/25/09, 3/12/10

In short, there are a number of issues in society that depend upon an understanding of what it means to use evidence.

Why should we care about the process?

Scientific debates are becoming more and more relevant to the world we live in. As information becomes easier and easier to access, it is critical that we understand how to use evidence to prove/disprove theories.

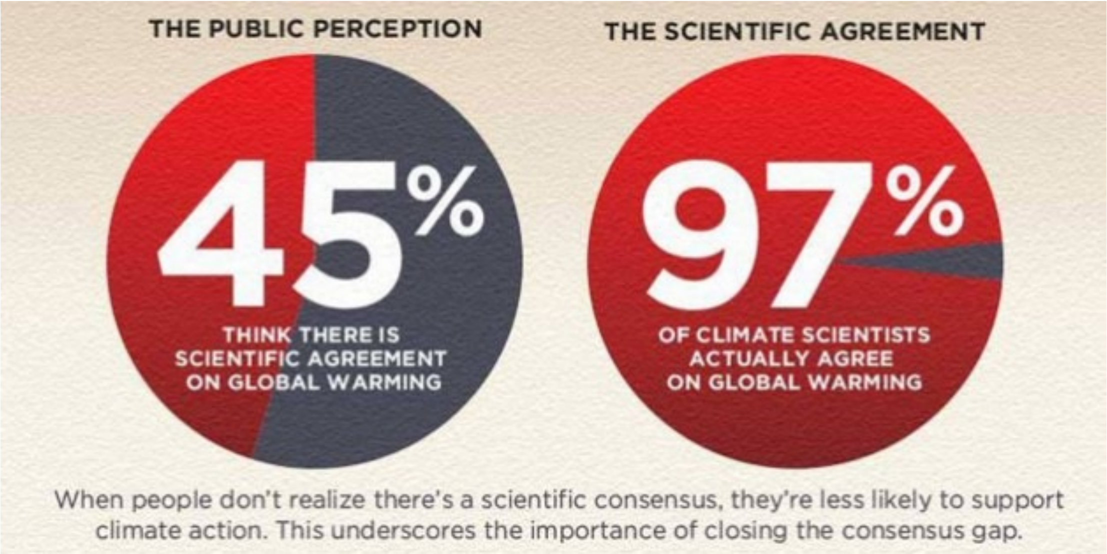

In short, there are a number of issues in society that depend upon an understanding of what it means to use evidence.

It is all about changing your mind!

We all have beliefs about how the universe works.

Science gives you a set of rules for figuring the universe out, and most importantly, for **changing your mind** when you encounter new evidence.

The rules of science say that if you believe something, you should be able to state exactly what evidence you have for your belief AND what evidence you would need to see to change your belief!

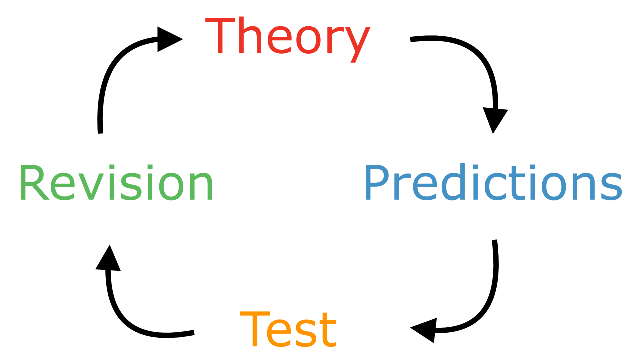

Then you can use the processes like **falsification** and **confirmation/ Bayes Theorem** to update your beliefs - which we will discuss next!

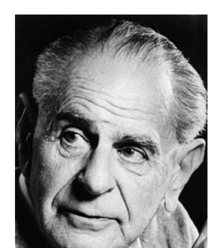

Karl Popper
1902-1994

Thomas Bayes
1701-1761

The challenge of infinite theories and finite evidence:

The Problem of Confirmation

Positive Evidence and The Problem of Confirmation

Recall that we learned that positive evidence will not guarantee that we learn the correct infinite set from a finite subset. This is because any given finite subset is compatible with multiple infinite sets.

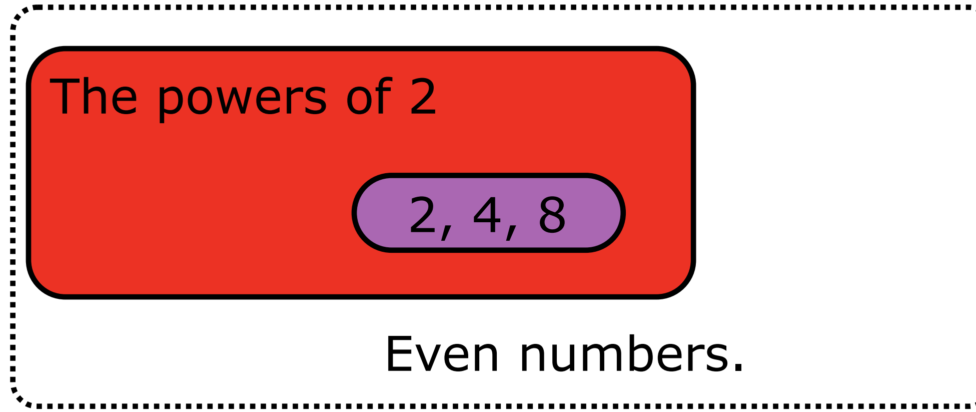

The same problem transfers to the infinite sets defined by scientific theories.

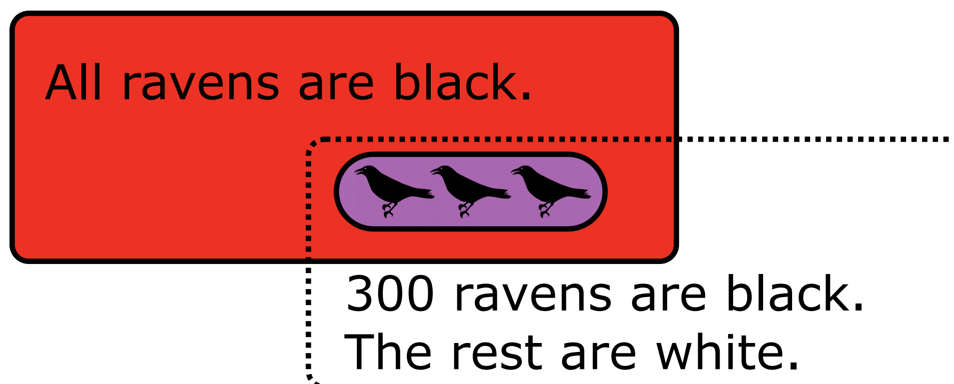

For example, let's say you've observed 300 black ravens. That finite subset is compatible with the theory that all ravens are black. It is also compatible with the theory that 300 are black, and the rest are white.

In fact, it is compatible with every theory that has at least 300 black ravens! (301, 302, 303…)

The Problem of Confirmation

We use the term **confirmation** to describe the process of finding evidence that supports a theory. In other words, observing examples that match your theory. In yet other words, finding positive evidence.

Although confirmation sounds like a good thing, it comes with a problem called **the problem of confirmation**. The problem of confirmation is that any given piece of positive evidence can be used to confirm (i.e., is compatible with) an infinite number of theories!

In other words, you can never confirm a single theory. Whenever you have positive evidence, it is evidence for an infinite number of theories.

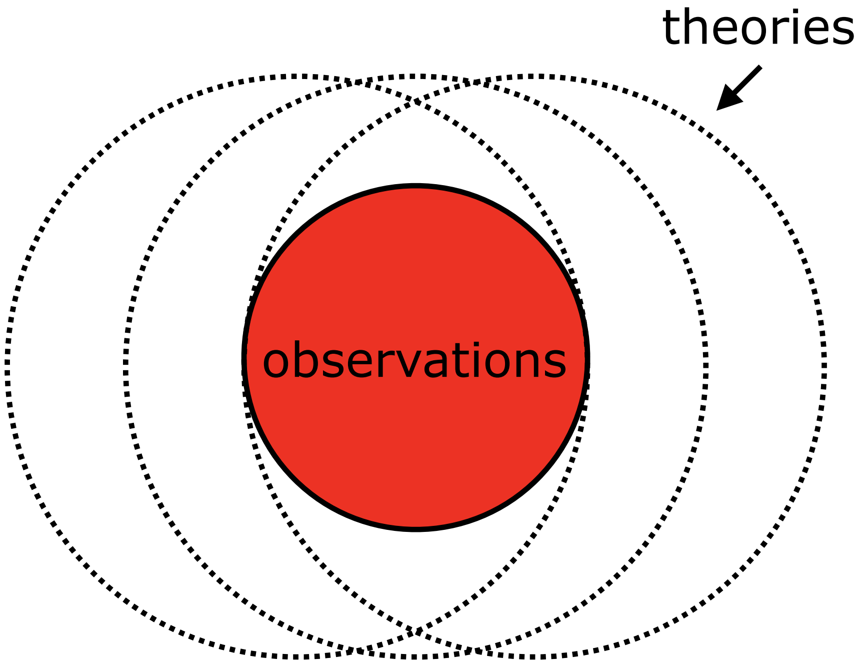

And that means you haven't really done anything at all.

One solution to the problem of confirmation: Falsification

Negative Evidence and Falsification

Recall that we learned that negative evidence, what is NOT in the set, can help us.

We saw that we could use it strategically to test different theories…

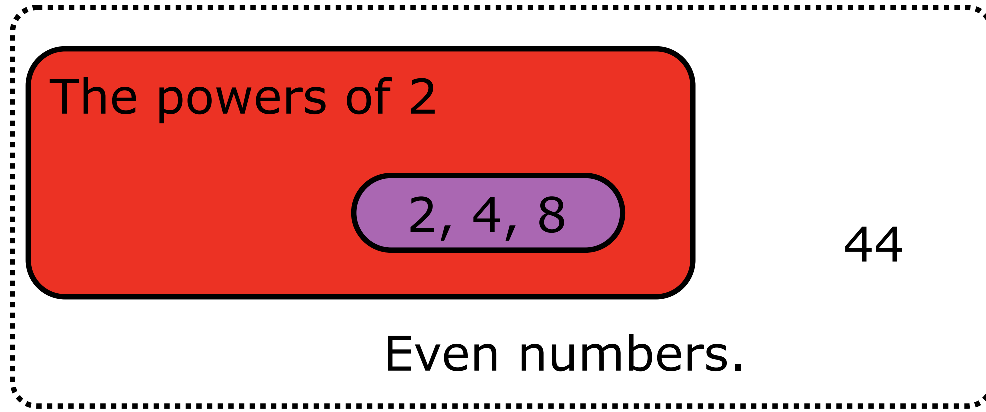

We can do the same thing with scientific theories.

If the theory is true, all ravens should be black. We can look for what should not be in the set - white ravens.

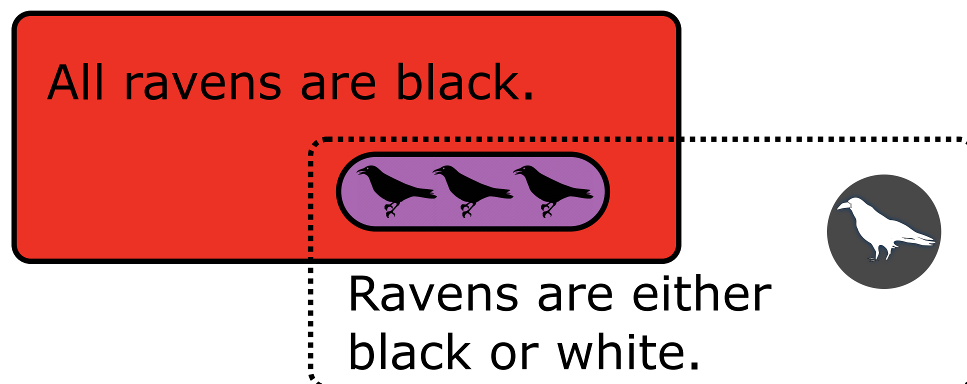

If we find a white raven, the theory that all ravens is black is **not true**. We say that we have **falsified** the theory that all ravens are black. And then we should only consider theories that allow for both black and white ravens.

Falsification in science

The philosopher Karl Popper proposed **falsification** as a way around the problem of confirmation.

falsification: The process of attempting to prove a theory wrong (falsify it).

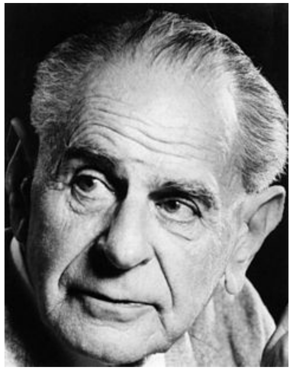

Karl Popper
1902-1994

For Popper (and other believers in falsification), a theory is only scientific it can be falsified. In other words, scientific theories have to take risks. They have to make predictions that could potentially prove them wrong. If a theory can't be falsified it is not a scientific theory!

This is a powerful idea, because it constraints what counts as a scientific theory. **You must be able to say what evidence would disprove the theory**. If you can't do that, it is not testable, and therefore not scientific.

The limits of falsification

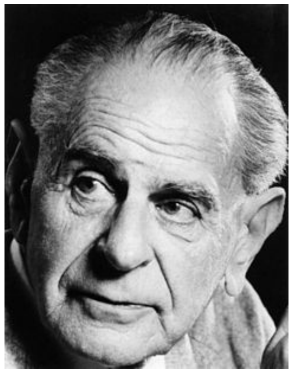

Karl Popper
1902-1994

One of the most interesting aspects of Popper's theory of falsification is that it completely denies the existence of confirmation. Falsification says "Confirmation is a myth." You cannot support a theory with evidence.

If a prediction is shown to be false, then the theory is falsified.

If a prediction is shown to be true, then **we can't say anything** about the theory. All we can say is that the theory has not yet been falsified.

In other words, under falsification, you **can never prove a scientific theory to be true**. You can only say that it hasn't been disproven.

But humans like confirmation…

Many people find it odd to say that you can't prove a scientific theory. In fact, many people have the intuition that confirmation is real!

Here is a thought experiment to demonstrate that we tend to believe in confirmation:

Let's say that you are asked to build a new bridge. You have two choices of design:

1. An old design that has been used for hundreds of bridges, none of which have collapsed.

2. A brand new design that has never been tested before.

Which would you choose? Falsification says that both bridges are equal, because neither has been falsified yet. But I bet you'd prefer to drive on the one that has been tested hundreds of times! That is confirmation.

Putting confirmation back into science: Probabilities and Bayes Theorem

Probabilities may allow for confirmation

The problem of confirmation teaches us that positive evidence is compatible with an infinite number of theories.

But this does not mean that the evidence is equally compatible with each theory.

Let’s say you’ve observed 300 black ravens, and no white ravens.

 x 300

 x 0

This observation is compatible with an infinite number of theories:

300 ravens are black, the rest are white.

301 ravens are black, the rest are white.

302 ravens are black, the rest are white.

…

95% of ravens are black, 5% are white.

100% of ravens are black.

But these are relatively unlikely. It is unlikely that you just happened to find all of the black ravens and none of the white ones!

These are more likely.

The probability of a theory

This intuition suggests that, even though positive evidence is compatible with an infinite number of theories, positive evidence can suggest that some theories are more likely than other theories.

So what we want to do is develop a precise way to conclude how likely a theory is given a piece of positive evidence.

And here is an equation that might do it for us. It is called **Bayes Theorem**.

$$P(\text{theory} \mid \text{evidence}) = \frac{P(\text{evidence} \mid \text{theory}) \times P(\text{theory})}{P(\text{evidence})}$$

Don't worry. I will explain how this works over the next few slides. You also don't have to memorize this equation for the exam. I just want to show this to you because it is an incredibly important equation in science (and cognitive science), so it is something you should know about. On the exam itself, all you need to know is what Bayes Theorem does for us; not the equation itself.

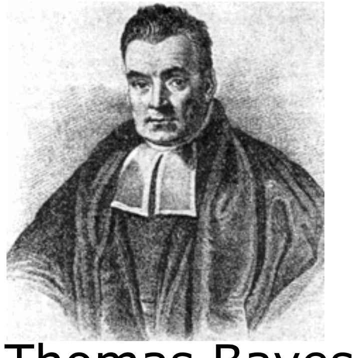

Thomas Bayes
1701-1761

Some basic terms

Probability: A mathematical statement about how likely an event is to occur. It takes a value between 0 and 1, where 0 means the event will never occur, and 1 means the event is certain to occur. (You can also think of it as a percentage 0% to 100%)

We write it like this: $p(\text{theory}) = .5$

Conditional Probability: The probability of an event **given that** another event has occurred.

We write it like this: $p(\text{theory}|\text{evidence}) = .9$

This says that the probability of our theory being true given the evidence that we observed is .9

This is what we care about - we want to know how likely our theory is given that we collected some evidence for it!

Getting a feel for conditional probabilities

What is the probability of being a dark wizard given that you are in slytherin?

p(dark wizard | slytherin) =

$$\frac{\text{number of dark wizards from Slytherin}}{\text{number of students from slytherin}} = \frac{30?}{5{,}000?} = \textbf{fairly low!}$$

Notice that if you flip the conditional around, the numerator is the same, but the denominator changes. This changes the probability entirely!

What is the probability of being from Slytherin given that you are a dark wizard?

p(slytherin | dark wizard) =

$$\frac{\text{number of dark wizards from Slytherin}}{\text{number of dark wizards}} = \frac{30?}{30?} = \textbf{very high!}$$

Bayes Theorem tells us how to calculate one conditional from its inverse!

Notice that this is just an equation. It has one condition probability on the left hand side. And it has the inverse (flipped order) on the right hand side, plus two basic probabilities.

$$P(\text{theory} \mid \text{evidence}) = \frac{P(\text{evidence} \mid \text{theory}) \times P(\text{theory})}{P(\text{evidence})}$$

Equations just tell us how to calculate things. This equation says we can calculate the thing on left if we know all of the numbers on the right. Notice that the thing on the left is really important — it is confirmation in science! So this is an important equation for us!

Actually applying Bayes Theorem to a theory is not always easy. It can be very difficult to calculate all of the probabilities on the right hand side. But nobody said science was easy!

A fun example of Bayes Theorem in action

The Monty Hall problem

Here are three doors. One has money behind it. The other two have goats.

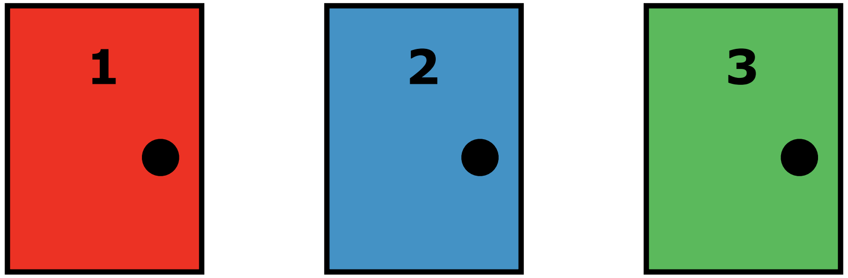

In the 1970s, there was a game show hosted by Monty Hall called Let's Make a Deal. In one of the games on the show, he showed contestants three doors, and asked them to choose one:

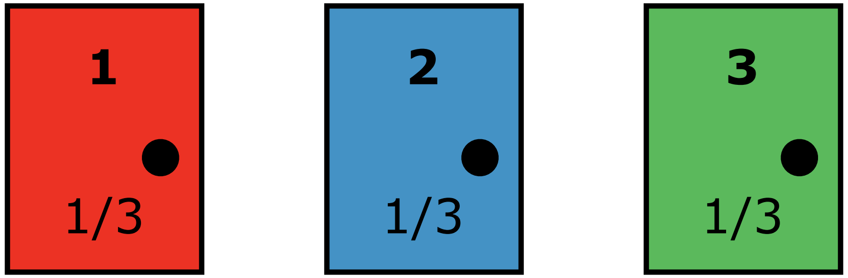

At the start of the game, the probability for each door that the money was behind it was 1/3.

The Monty Hall problem

To make things concrete, let's select door 1:

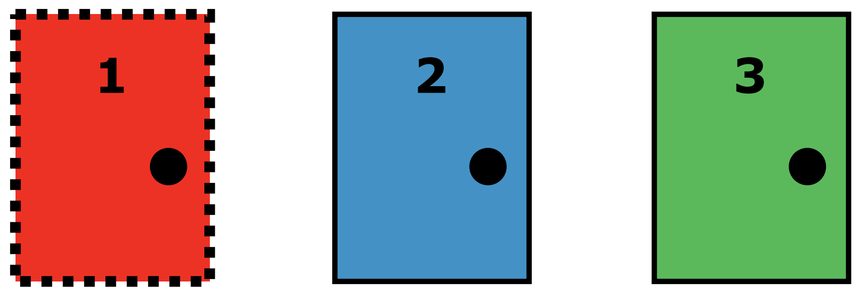

After contestants selected their door, Monty Hall would then select one of the other doors himself, and open it to reveal a goat:

Monty Hall would then ask the contestants if they wanted to keep their first choice, or switch to the remaining unopened door. **What would you do?**

This is about probabilities

In order to make a choice, you have to figure out the probabilities that the money is behind each door.

The reason this is tricky is that most people assume that nothing has changed in the probabilities. They were equal before the goat, and they are equal after:

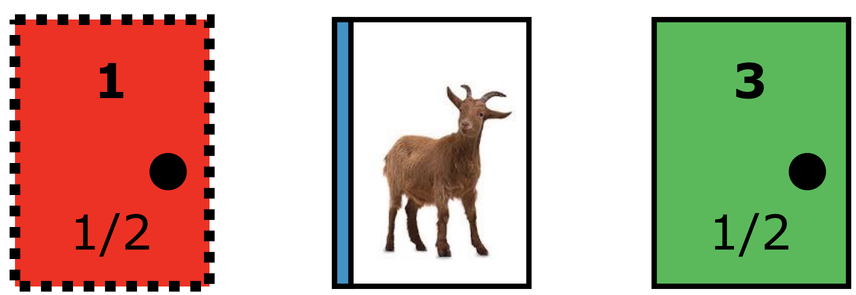

Each door was 1/3 before the goat. After the goat, there are only two options left, and they must sum to 1, so they must be 1/2 each.

But this is wrong!

In fact, the probabilities have changed:

The intuitive reason behind this is that Monty Hall knew which door hid the money. When he made his choice, he used this knowledge - he made sure that he did not open a door that hid money. So his action is evidence!

And once we see it as evidence, we can use Bayes Theorem to calculate exactly what the probabilities are for each door:

$$P(\text{prize } 1 \mid \text{open } 2) = \frac{P(\text{open } 2 \mid \text{prize } 1) \times P(\text{prize } 1)}{P(\text{open } 2)}$$

$$P(\text{prize } 3 \mid \text{open } 2) = \frac{P(\text{open } 2 \mid \text{prize } 3) \times P(\text{prize } 3)}{P(\text{open } 2)}$$

The first conditional probability

$$P(\text{prize } 1 \mid \text{open } 2) = \frac{P(\text{open } 2 \mid \text{prize } 1) \times P(\text{prize } 1)}{P(\text{open } 2)}$$

$P(\text{prize } 1)$ — This is 1/3. There are 3 doors, and the TV show could choose any of them. The starting probability (prior probability) for each door is 1/3.

$P(\text{open } 2 \mid \text{prize } 1)$ — This is 1/2. If the prize is behind door 1, then Monty can choose either door 2 or door 3.

$P(\text{open } 2)$ — This is the tricky one. The answer is 1/2, but I need the entire next slide to show you how we get that.

Calculating p(open 2)

The tricky bit for calculating the evidence is that you have to consider every possible theory (prize 1, prize 2, prize 3), and calculate the probability of the data (open 2) under each theory.

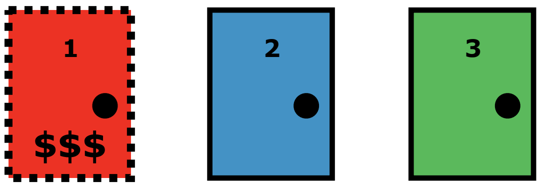

=1/2

If the prize is behind door 1, the probability of opening 2 is 1/2.

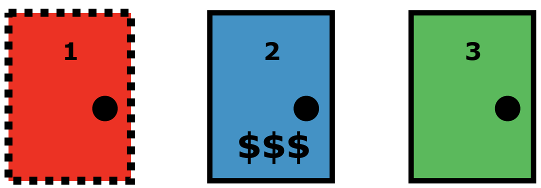

=0

If the prize is behind door 2, the probability of opening 2 is 0. Monty can't open that door.

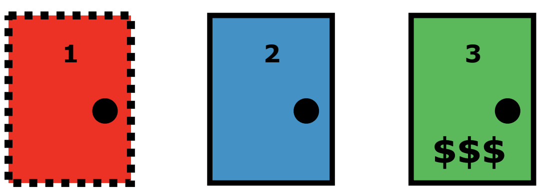

=1

If the prize is behind door 3, the probability of opening 2 is 1. Monty can't open door 1 because the contestant chose it. He can't open 3 because it has the prize. So 2 is the only one.

There are 3 theories, each with a prior probability of 1/3. We weight each theory by its prior probability, which means multiplying each by 1/3, and then sum. This tells us how likely open 2 would be out of all possible worlds:

$$P(\text{open } 2) = (1/3 \times 1/2) + (1/3 \times 0) + (1/3 \times 1) = 1/2$$

The first conditional probability

$$P(\text{prize } 1 \mid \text{open } 2) = \frac{P(\text{open } 2 \mid \text{prize } 1) \times P(\text{prize } 1)}{P(\text{open } 2)}$$

$$P(\text{prize } 1 \mid \text{open } 2) = \frac{1/2 \times 1/3}{1/2}$$

$$P(\text{prize } 1 \mid \text{open } 2) = 1/3$$

We can actually take a shortcut now. Since the prize has to be behind door 1 or door 3, and we know door 1 is 1/3, and we know probability must equal 1, then that means that probability for door 3 must be 2/3! But let's do the calculation anyway.

The second conditional probability

$$P(\text{prize } 3 \mid \text{open } 2) = \frac{P(\text{open } 2 \mid \text{prize } 3) \times P(\text{prize } 3)}{P(\text{open } 2)}$$

P(prize 3)

This is 1/3. There are 3 doors, and the TV show could choose any of them. The starting probability (prior probability) for each door is 1/3.

P(open 2 | prize 3)

This is 1. Monty can't choose 1 because the contestant chose it. He can't choose 3 because it has the prize. He has to choose 2

P(open 2)

We already know that this is 1/2 by the previous calculation.

The second conditional probability

$$P(\text{prize 3} \mid \text{open 2}) = \frac{P(\text{open 2} \mid \text{prize 3}) \times P(\text{prize 3})}{P(\text{open 2})}$$

$$P(\text{prize 3} \mid \text{open 2}) = \frac{1 \times 1/3}{1/2}$$

$$P(\text{prize 3} \mid \text{open 2}) = 2/3$$

This is exactly what we calculated with our shortcut. So we can see that Bayes Theorem really works.

Bayes Theorem is not intuitive!

What we thought:

What Bayes gave us:

We can show that Bayes gives us the correct result with simulations:

https://www.rossmanchance.com/applets/MontyHall/Monty04.html

Lesson 1: Bayes theorem works. We can see that the math and the simulations match up.

Lesson 2: Our gut reactions are not always correct. The Philosophy of Science (including falsification and Bayes) gives us tools to make sure that we are updating our beliefs correctly. It is important that we apply them rather than just going with our gut reactions.

Conclusions

A theory is only scientific if it makes **testable predictions**. We can then go out and test those predictions to see if the theory is correct or not. If it is not, we can revise the theory, and start the cycle over again.

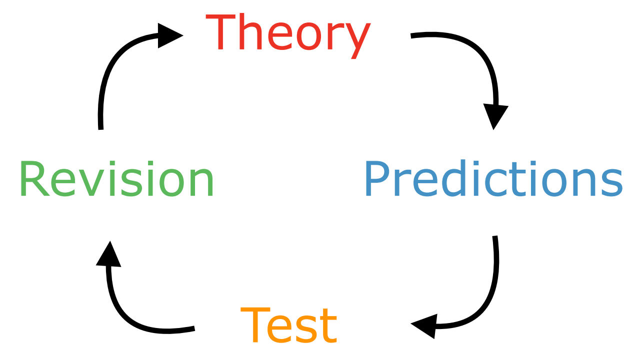

Confirmation is the process of finding evidence that supports a theory (positive evidence).

> **The problem of confirmation** is that finite observations are compatible with an infinite number of scientific theories.
>
> We can use **Bayes Theorem** to partially overcome this, by deriving probabilities for theories from positive evidence.

Falsification is the process of finding evidence that disproves a theory with certainty.

> Falsification is the gold standard of science, but there are circumstances where it isn't possible.