T. Parent

# INTRO TO PROBABILITY

## ***1. The Basics***

Why care about probability? Well, recall that the worth of a non-deductive argument largely depends on how *probable* the conclusion is, given the truth of the premises. Since the notion of probability is key with non-deductive arguments, logicians look to the mathematics of probability to calculate in an exact way the probability of a conclusion, in light of its premises.

Hereafter, when we speak of the probability of some statement **P**, we will write "*p*(**P**)." Note that **P** will be a statement*, not a number!!* But crucially, "*p*(**P**)" IS a number; it is the *probability* of the statement being true. [Thus, "*p*( )" can be seen as a function that takes a statement as input, and outputs a number that represents its likelihood.]

Example: *Given a fair die, what's the probability that you will roll a four? That is, what number is* p*(Roll Four)?*

Since there are six faces to a die, and only one face shows four, then there's one chance in six that you'll roll a 4. That is:

$$p(\text{Roll Four}) = 1/6.$$

As in this example, a probability is always a number between 0 and 1 (inclusive):

$$0 \leq p(\mathbf{P}) \leq 1$$
$$0 \leq p(\mathbf{P}/\mathbf{Q}) \leq 1$$

Note that there are two claims here, where the second claim is about (what we may call) the *posterior* probability of **P**. The posterior probability of **P** is relevant when we want to learn not the probability of **P** itself, but rather the probability of **P** *once we have obtained some evidence* **Q**. That is, sometimes we want to *update* the probability of **P**, after we have gained some new information **Q**. Thus, it is the probability of **P** *after* ("posterior") we have learned of **Q**, and it is written symbolically as "*p*(**P**/**Q**)."

<u>Example</u>. *Given a fair die, what's the probability of rolling a four, given an even number? That is, what is the posterior probability of rolling a four, once we know an even number was rolled?*

Since there are 3 even numbers on a die (two, four, and six), then if we are *given* that an even number was rolled, the chances of it being four specifically is 1 in 3. That is to say:

*p*(Roll Four/Even) = 1/3.

Note that when we talk about posterior probabilities, sometimes we refer to the initial probability as the *prior* probability of **P** (since it is before the evidence comes in). Yet the prior probability of **P** is just the same thing as *p*(**P**). Thus, in our example, the prior probability of rolling a four is 1/6, since that just is the probability of rolling a four (absent any additional information).

The claim that probabilities are between 1 and 0 constitutes the first "axiom" of probability. The second axiom is the claim that if **P** cannot be false, then *p*(**P**) = 1. This enables us to infer the probability of ~*p*, given the probability of *p*, as per the following rule:

**[NOT Rule] *p*(~P) = 1 – *p*(P)**

The thought here is that, since **P** ∨ ~**P** cannot be false, then by Axiom 2, *p*(**P** ∨ ~**P**) = 1. So to know the probability of ~**P** *specifically*, we can subtract *p*(**P**) from 1.

## ***<u>2. Joint Probabilities and the Gambler's Fallacy</u>***

If I have a fair coin, then *p*(Heads) = 50%. But suppose I am planning to toss the coin twice and want to know the probability of getting heads on both tosses. This means discerning the probability of *two* statements being true: I get a Heads on Toss 1, and I get a Heads on Toss 2. This can be represented as: *p*(HeadsT1 & HeadsT2).

In figuring this out, note that in 50% of cases, we will get Heads on the first toss, and in 50% of *those* cases, we will also get a Heads on the second toss. So basically, in 50% of 50% of cases, we will get two heads. Equivalently: In 25% of all cases, we will get two heads. This thinking can be generalized into a formula that applies to any two statements **P** and **Q**:

**[Simple AND Rule] If Q is independent of P, then *p*(P & Q) = *p*(P) × *p*(Q)**

So in our example, this formula tells us that *p*(HeadsT1 & HeadsT2) = .5 x .5 = .25.

But it is crucial that this reasoning is sound only because HeadsT2 is statistically *independent* of HeadsT1: We know that 50% of 50% of cases will yield two tails, partly because we know that getting a Heads on Toss 1 *does not affect the probability* of getting Heads on Toss 2.

In contrast, however, suppose that I have a “magic coin” where getting Heads on one toss decreases the chances of getting Heads on the next toss. To make things simpler, suppose that the magic coin has never been tossed before and *p*(HeadsT1) = 50%. But the magic coin is such that, assuming I get Heads on Toss 1, it is then only 5% likely that I will get Heads on Toss 2. That is, *p*(HeadsT2/HeadsT1) = .05. Then, I should reason as follows. 50% of the time I will get Heads on Toss 1, and in those cases, there is then only a 5% chance of HeadsT2. So basically, I have 5% chances of getting two heads *only half of the time*. Equivalently: I have only 2.5% chances of getting two heads. My reasoning in this sort of case—when the second event may depend on the first—is generalized in the following rule:

**[AND Rule] *p*(P & Q) = *p*(P) × *p*(Q/P)**

This is just like the [Simple AND Rule] except that *p*(**Q**) has been replaced by *p*(**Q**/**P**). That is quite apt. With the magic coin, assuming I get Heads once, the chances of Heads a second time is not 50%, as it would be with a normal coin, but only 5% given HeadsT1. And so, the probability of getting Heads on both tosses is .5 x .05 = .025.

This also reveals what it means for **Q** to be “independent” of **P**. It is a matter of whether the posterior probability of **Q**, given **P**, us the same as the prior probability of **Q**:

Definition: **Q** is (statistically) *independent* of **P** iff $p(\mathbf{Q}) = p(\mathbf{Q}/\mathbf{P})$

Intuitively, "**Q** is independent of **P**" means that the probability of **Q** is unaffected by the truth or falsity of **P**; its probability *doesn't care* what **P** is doing.

WARNING: *Order matters.* Even if **Q** is dependent on **P**, that does NOT mean **P** is dependent on **Q**! Think of the magic coin: Even though the second toss depends on the first, the first does not depend on the second. The chances of HeadsT1 is 50%, regardless of whether HeadsT2 occurs.

One reason why all this is important is in helping us avoid the *gambler's fallacy*. There is a tendency to think that later events depend on earlier events, even when they do not. A classic example is when a gambler basically treats a coin as "magic," believing that a Heads on three successive tosses makes it better to bet Tails on the next toss. Or, more commonly, a gambler will regard a roulette wheel as "magic" in that, if the wheel comes up Red on three successive spins, then it is better to bet Black for the next spin. But really, coin tosses and roulette spins are independent of each other (assuming that nothing is rigged). On a fair coin, the chance of Heads is always 50%, even if you have just witnessed 100 tosses where it came up Heads.[1] Gambling examples make the point obvious, but less obvious examples occur, e.g., when a doctor mistakenly assumes that a patient's risk for diabetes is independent of economic class.[2] Keep an eye out for mistaken assumptions about what is independent of what.

## ***3. Disjoint Probabilities and the Error of Double Counting***

Sometimes we want to know not the probability of two events occurring, but rather the probability of *at least one* of two events occurring. For example, we might want to know what the probability is of getting either an A or A- in this class last semester. Suppose here that there were 50 students in the class last semester, and 3 got an A, whereas 4 got an A-.

---

[1] Of course, if you witness 100 Heads in a row, that seems like evidence that the coin is *not* fair. After all, the prior probability of getting 100 Heads in a row is quite small. But the point above is that *if* a coin is fair, then 100 Heads in a row does not make Heads any less likely on the 101st toss.

[2] You may find this surprising, but many health factors are statistically dependent on socio-economic factors like race, class, and gender. Public health professionals thus commonly speak of the "social determinants of health." For more information, see https://health.gov/healthypeople/priority-areas/social-determinants-health.

Then if you pick one of the students randomly, the chances that you will get an A student is 3/50…and the chances you will get an A- student is 4/50. So the chances that you will get either type of student is 7/50. The thinking here generalizes into the following rule:

**[Simple OR Rule]   If P and Q are mutually exclusive, then $p(P \lor Q) = p(P) + p(Q)$**

Like the [Simple AND Rule], this rule works only under a certain sort of assumption. In this case, it is that **P** and **Q** *mutually exclude* each other, meaning that you cannot have *both* **P** and **Q** being true. That applies to our example with the students; after all, a student cannot get both an A and an A- in the course. More broadly:

Definition: **P** and **Q** are *mutually exclusive* iff it cannot be true that **P** & **Q**.

Notice that getting an A and an A- are mutually exclusive, even though many students earned neither an A nor an A-. So to say that **P** and **Q** are mutually exclusive does *not* imply that *at least one* must be true. It is rather to say that *at most one* can be true.

Now, we often deal with statements that are not mutually exclusive. Consider the chances of a student getting an A on the first exam or an A on the second exam. Suppose that 5 out of 50 got an A on the first exam, whereas 10 out of 50 got an A on the second exam. Importantly, it does not follow that 15 out of 50 students got an A on either the first or second exam. That's because some of the students who got an A on the first exam might have *also* gotten an A on the second exam. Thus, if one student got an A on both exams, then really, only 14 students would have gotten an A on the first or second exam. If we said instead that 15 students got an A on either exam, we would be "double counting" the one student who got an A on both exams. We would be mistakenly treating the one student as if s/he were two students.

Thus, if someone were to mistakenly apply the [Simple OR Rule] to the case, they would be thinking that the chances of getting an A on either exam is 5/50 + 10/50 = 15/50 = 30%. But since only 14/50 got an A on either exam, the probability here is actually lower.

So when we are dealing with events that are not mutually exclusive, we need to subtract any "overlap" so to avoid double counting. The general rule here is this:

**[OR Rule]** $p(P \vee Q) = p(P) + p(Q) - p(P \,\&\, Q)$.

Thus, in our example about the exams, this rule tells us that the probability of getting an A on either exam is 5/50 + 10/50 – 1/50 = 14/50 = 28%.[3]

## *4. Calculating Posterior Probabilities*

So far, we have been dealing with cases where the posterior probabilities are fairly obvious. But very often that is not so. Fortunately, however, there are formulae which allow us to calculate what the posterior probability is, using other information. Here is one:

**[Bayes' Theorem "BT"]** **If $p(Q) \neq 0$, then:**

$$p(P/Q) = \frac{p(P) \times p(Q/P)}{p(Q)}$$

Understanding the rationale for this formula is easiest if we consider an example.

Example: *Suppose I draw a card from an ordinary deck of 52 playing cards, and I do not draw a Spade. What's the probability that I drew the Queen of Hearts?*

The prior probability of drawing the Queen of Hearts is 1/52, and the prior probability of drawing a non-Spade is 3/4. The only other number we need to use BT here is the probability of drawing a non-Spade, given that I drew the Queen of Hearts. That probability is 1—after all, if we are *given* that I drew the Queen of Hearts, then I cannot have drawn a Spade! So plugging in these numbers:

$$\frac{p(\text{QHearts}) \times p(\sim\text{Spade}/\text{QHearts})}{p(\sim\text{Spade})} = \frac{1/52 \times 1}{3/4} = \frac{1}{52} \times \frac{4}{3} = \frac{4}{156} = \mathbf{\frac{1}{39}}$$

[3] In saying that 1/50 is the probability of getting an A on both exams, we were NOT assuming that those events were independent. (This is as it should be— getting an A on exam 1 may increase the chances of getting an A on exam 2.) After all, if we apply the [AND Rule] to the case, we get the same 1/50 answer. We need only be clear that the probability of getting an A on exam 2, *given* an A on exam 1, is 1 out of 5 students. (Only 5 students got an A on exam 1, and exactly one of them also got an A on exam 2.) Then, the [AND Rule] tells us $p(\text{AExam1} \,\&\, \text{AExam2}) = p(\text{AExam1}) \times p(\text{AExam2}/\text{AExam1}) = 5/50 \times 1/5 = 5/250 = 1/50$.

You might have been able to figure this out, even without applying BT. You might reason that, if I did not draw a Spade, then since there are 13 Spades out of 52 cards, I must have drawn one of the other 39 cards. And since exactly one of those 39 cards is the Queen of Hearts, it must be that there's a 1/39 chance of getting the Queen of Hearts, assuming a non-Spade was drawn.

Your reasoning here is essentially what BT captures. BT says you can figure out what *p*(**P**/**Q**) is, if you pretend like the only possibilities that exist are the possibilities where **Q** is true. Very roughly, you figure out what fraction of those **Q**-possibilities are ones where **P** is also true. That fraction is the probability of **P**, given **Q**. (Things are a little more complicated when **Q** given **P** is uncertain—this is why the numerator in BT is not just *p*(**P**), but rather the product of *p*(**P**) and *p*(**Q**/**P**).)

Importantly, we can extend this exercise to determine the likelihood of the conclusion in one sort of argument. Consider specifically the following non-deductive argument:

> (P1) 1/52 cards is the Queen of Hearts, and 39/52 cards are not Spades.
> (P2) I drew a non-Spade.
> (C) So, I did *not* draw the Queen of Hearts.

Assuming the premises, we saw that I had a 1/39 chance of drawing the Queen of Hearts. So by the [NOT Rule], (C) has a 38/39 chance of being true. Since that is significantly above 50%, it is significantly more likely that the conclusion of this argument is correct. To that extent, the argument is a good one. Granted, there remains approximately 2.6% chances that the conclusion is wrong. So it is completely possible that (C) is false (which is to be expected in a non-deductive argument). But it is entirely more likely that it is true.

Nonetheless, you might think it is irrational to conclude (C), even assuming the truth of the premises. But since non-deductive arguments always carry risk, the risk *per se* does not make the conclusion irrational. Certainly, it is better to conclude that (C) is true rather than *false*. But in some contexts, it may even be better to conclude that (C) is true rather than claim ignorance. After all, a 2.6% chance of error is relatively small—and depending on what's at stake, it may be better to take the risk rather than just shrug your shoulders.

In any case, this illustrates how we can precisely judge the worth of a non-deductive argument using probability calculations.

## ***5. The Base-Rate Fallacy***

Human beings often have a terrible time intuiting probabilities. Because of our limitations, the notorious *base-rate fallacy* often occurs when estimating posterior probabilities. The following is an example to illustrate.

Example: *You have been called to jury duty in a town where there are two taxi companies, Green Cabs Ltd. and Blue Taxi Inc. Blue Taxi uses cars painted blue; Green Cabs uses green cars. Green Cabs Ltd. dominates the market, with 85% of the taxis on the road.*

*On a misty night, a taxi sideswiped another car and drove off. A witness says it was a blue cab. The witness is tested under conditions like those on the night of the accident, and in 80% of cases, she correctly reports the color of the cab. That is, regardless of whether she is shown a blue or a green cab in misty evening light, she gets the color right 80% of the time.*

*You conclude on the basis of this information:*

(a) *The probability that the sideswiper was blue is 80%*
(b) *It is more likely that the sideswiper was blue, but it is slightly less than 80%.*
(c) *It is just as probable that the sideswiper was green as that it was blue.*
(d) *It is more likely than not that the sideswiper was green.*

Pick an answer before reading on.

This question was devised by psychologists Amos Tversky and Daniel Kahneman. After extensive testing, they found that many people think that (a) or (b) is correct. Very few people think that (d) is correct. Yet (d) is the right answer!

In this case, people are prone to ignore *the probability that a randomly selected taxi will be green vs. blue.* This is the prior probability; it is known prior to the additional evidence (given by the witness). The prior probability here is also known as the *base rate*: It is the rate at which any given cab will be green vs. blue.

The fallacy consists in ignoring the base rate when answering the question. The base rate effectively means that the witness says "blue" incorrectly more often than she says "green" incorrectly. And since she answered "blue," the numbers work out so that a green cab remains more likely.

In more detail: Assume for simplicity's sake that there are 100 cabs total, so that there are 85 green cabs and 15 blue cabs. Since the witness is 80% reliable, we'd expect the following:

1. She would correctly identify 80% of the green cabs, and 80% of the blue cabs. This would mean she would correctly identify 68 greens and 12 blues.
2. Conversely, it means she would misidentify 17 greens as "blue" and 3 blues as "green."

Thus, she would identify 29 cabs in total as "blue," but only 12 would be blue. So given that the witness says "blue," the chances of her being correct are 12/29, which is about 41%. It thus remains less than 50% likely that the cab was blue.

But instead of thinking through the case in this manner, you can simply plug the numbers into Bayes' Theorem and get the same 12/29 answer. This is left as a homework exercise.

Here is an even more dramatic example to show why the base rate matters. (This example was mentioned in Forseman et al., ch. 6, but it's worth reviewing in more detail.)

<u>Example</u>: *There is a very rare and very terrible disease, "fictionitis." In the general population, only 1 person in 10,000 has the disease, but you decide to get tested anyway on a whim. The test is 99% accurate: If you have the disease, the test says YES with probability 99%—and if you do not have the disease, it says NO with probability 99%.*

*You get your results back and the test says YES. What are the chances that you have the disease?*

*(a) 99%*

*(b) Slightly less than 99%*

*(c) 50%*

*(d) About 1%*

The answer in this case is also (d)—but the trick is to see *why* that is the right answer. The rough idea is this: The chances of having the disease is .01% for any member of the population, hence, .01% for any person that is tested. In contrast, the test gives the wrong answer to 1% of those people. So roughly, the likelihood of disease is still *100 times less* than the chances of a correct YES answer.

In more detail: If a total of 1,000,000 people are tested, then only 100 will have the disease. This represents the prior probability of having the disease. The reliability of the test then means that 99 of these 100 will get a YES result.

But what about the other 999,900 people who don't have the disease? The reliability of the test here means that 1% of 999,900 will get an incorrect YES result. But that ends up being *9999 people*!

So what are the chances that *your* YES result is correct? Well, there are 99+ 9999 = 10,098 people who got a YES result. But 9999 of 10,098 had an *incorrect* YES result. So the chances that your YES result is incorrect is 9999/10,098, which is approximately 99%!

Thus, even though the test is 99% reliable, it is still 99% likely that your YES result is incorrect. Since the base rate of the disease is *so* low, it ends up making it much more likely that your test result was a false positive, despite the high reliability of the test.