Bayesian Thinking: The Rule and the Consequences

QTM 110

KEVIN MCALISTER

Administrative Stuff

Problem Set #4 will be posted on Thursday

- Sorry for the back to back problem sets!
- Want to make sure you have plenty of practice!
- One question on IV designs and one for Bayes' Theorem

Midterm #2 is two weeks from today:

- Will cover everything since the last midterm (+ Bayesian thinking) and experiments generally
- More details next week

Administrative Stuff

We're in November, so it's time to start thinking about final exam times:

- December 11th 11:30 – 2:00 for Section 1
- December 16th 3:30 – 5:00 for Section 2

Some requests to reschedule:

- Will be offering an additional exam time on December 12th 12:00 – 2:30. Can take with the other section (space permitting).
- Requests to take the exam at a time other than your assigned time should be made using the survey on the Canvas site
- Please complete before November 25th to ensure that your alternate exam time can be accommodated!

Where do we go from here…

- We’ve spent a lot of time thinking about **causality**
  - How can we say that X causes Y
- One very important class of questions
- But there are others:
  - Questions of probability/counting/description
  - Questions of prediction
  - Questions about how much we can trust findings
- The last part of this class will cover some of these topics

People v. Collins

Los Angeles, 1964

An elderly woman was walking down an alley, pulling a basket of groceries, with her purse resting on top of the groceries.

Suddenly, she was pushed down to the ground from behind and her purse was stolen, but she didn’t get a good look at the perpetrator.

An eye witness saw a woman run out of the alley and enter a yellow car.

The witness also didn’t get a **good** look at anything.

He didn’t remember the make of the car,

People v. Collins

But he remembered that the woman running was white and had a ponytail.

He also saw that the driver of the car was black with a beard and mustache.

Based on this eye witness testimony, police later arrested Malcolm and Janet Collins and charged them with the robbery.

Malcolm was a black man with a beard and mustache.

Janet was a white woman with a pony tail.

And they drove a yellow car.

The “Math”

Prosecutors brought in a mathematician to testify regarding the chances that, **based on this evidence alone**, Malcolm and Janet were guilty of the robbery.

The mathematician concluded that there was only about a 1 in 12 million chance that the couple was innocent.

The testimony went as follows:

- If we just arrested an innocent couple at random, it’s very unlikely that
- The husband would be black with a beard and mustache
- The wife would be white with a pony tail,
- They would drive a yellow car.

The "Math"

Probabilities of various characteristics

- 10% chance black with beard
- 25% chance of a man with moustache
- 10% chance of white woman with pony tail
- 33% chance of white woman with blonde hair
- 10% chance of yellow car
- .1% chance of an interracial couple

Assuming independence: 1 in 12 million chance of a random couple meeting all of these criteria. Must have been them, right?

A jury found the couple guilty, and newspapers praised prosecutors for making such a quantitatively compelling case.

What problems do you see with this analysis?

What’s wrong with this analysis?

Basic Problem: These events are not independent.

- Having a beard is correlated with having a mustache
- Those may be correlated with driving a yellow car, etc.
- So you can’t just multiply the probabilities together to get the joint probability.

Account for the correlation in traits, we might get to a **1 in 500,000** chance that a randomly selected innocent couple would include a black man with a mustache and beard and a white woman with a ponytail who drive a yellow car.

- Larger than 1 in 12 million, but still a really low number— probably enough evidence for a jury to convict.

If we were statisticians, we’d call this a p-value – the probability that we would see this if the null (innocence) were true – P(Match | Innocent)

We’d say that we can strongly reject the null hypothesis of innocence, because it’s very unlikely that an innocent couple would have these characteristics.

What else has gone wrong?

Something much more important!

We asked the wrong question.

Grant that the eye witness was flawless

- The crime was perpetrated by a white woman with a pony tail who drive a yellow car and a black man with a beard and mustache

Grant that the two criminals are likely married or dating.

Given that information, we'd like to know

What's the probability that Malcolm and Janet are guilty, given that a crime was committed by a couple matching their physical characteristics?

The question we want versus the question we answered

THE QUESTION WE WANT THE ANSWER TO

What's the probability that Malcolm and Janet are guilty, given that a crime was committed by a couple matching their physical characteristics?

Or

Pr(innocent|matching characteristics)

THE QUESTION THAT WAS ANSWERED

What's the probability that a randomly selected innocent couple would match these characteristics,

Or

Pr(matching characteristics|innocent).

The answers to these two questions might be different. Let's see how.

The difference

Suppose we know the answer to the wrong question:

- Pr(matching characterstics|innocent) = 1 in 500000 or .000002

We want to compute **Pr(innocent|matching characteristics)**

What else do we need to know?

Bayes' theorem gives us the answer.

It tells us that

$$\Pr(\text{Innocent}|\text{Matching}) = \frac{\Pr(\text{Matching}|\text{Innocent}) * P(\text{Innocent})}{\Pr(\text{Matching})}$$

Let's derive this theorem from the ground up.

I Thought This Lecture Was On Bayes' Theorem?

- To talk about **Bayes' theorem** we have to learn about **conditional probabilities**
- To talk about **conditional probabilities** we have to talk about…

What is Probability?

- The **probability** of an event A is the chance that A occurs (or the relative frequency with which an event occurs)
  - Let's denote this P(A)
  - A number between 0 and 1, or 0% and 100%, inclusive
  - Some examples:
    - What is the probability that the sun will rise tomorrow?
    - What is the probability that I'll end class before 3:45 today?
    - What is the probability that a fair coin lands on heads?
    - What is the probability that tomorrow is a Friday?

What is Probability?

- The **probability** of an event A is the chance that A occurs
  - Let's denote this P(A)
  - A number between 0 and 1, or 0% and 100%, inclusive
  - Some examples:
    - What is the probability that the sun will rise tomorrow? Around 1.
    - What is the probability that I'll end class before 3:45 today? Around 0.
    - What is the probability that a fair coin lands on heads? .5
    - What is the probability that tomorrow is a Friday? 0. It just won't be.
  - Intuitive. Just formalizes something you already do every day.

What is Probability?

- Example: event A = I play a round of disc golf today

  - The probability that I play a round of disc golf today
    = P(A) = 0.25 or 25%

Probability Space

- Think of a set of possible outcomes, a **sample space S**
  - **S** could be {played disc golf, didn't play disc golf}
- The sum of the probabilities of all possible events in S has to equal 1, or 100%
  - *Something* has to happen
  - And the totality of probabilities must add up to 1

Sample Space S
Total Area = 1.0, or 100%

Independent Events

- Now let's consider two **independent events**, A and B
  - A = I played disc gold today; P(A) = 25%
  - B = my wife Dana wore a black sweater; P(B) = 20%
  - Events are unrelated; knowing if I played disc golf (the cold doesn't deter me) tells you nothing about whether Dana wore a black sweater (she's perpetually cold, even in the summer)

- P(I played disc golf *and* Dana wore a black sweater) =
  - **P(A and B) = P(A) x P(B)** = 0.25 x 0.2 = 0.05 or 5%
    - Works *only* for independent events

Marginal Probabilities

- Example: event A = I play disc golf today
  - The probability that I play disc golf = P(A) = 0.25 or 25%
  - Probability of one event in total, disregarding any other events = **marginal probability**

Conditional Probability

- Example: event A = I play disc golf
  - The probability that I play disc golf = P(A) = 0.25 or 25%

- Now add a second event B = it rained today
  - The probability that it rained today = P(B) = 0.20 or 20%

Conditional Probability

- Whether I play disc golf depends ("is **conditional** on") if it's raining
  - P(disc golf|rain) = 0.05 or 5%
    - "|" is called a **pipe**
    - Read as "the probability of playing disc golf given rain" or "the probability of playing disc golf conditional on rain"
  - This is a **conditional probability** = **P(A|B) =** Probability of event A given event B = **P(A and B)/P(B)**
  - Sidenote (for now): I play disc golf 30% of days when it *doesn't* rain

Joint Probability

- **Joint probability** is just the probability of two events, such as "Today I play disc golf *and* it's raining"

  - In words: 20% of the time it rains. 5% of *that* time I play disc golf
    - 5% of 20% = 1% of the time I play disc golf in the rain
  - In numbers: 5% x 20% = 0.05 x 0.2 = 0.01 or 1%
  - In probability:
    P(Disc golf *and* rain) = P(Disc Golf|rain)*P(rain) = 0.05*0.2 = 0.01 or 1%

  - More generally,
    Joint Probability = P(A *and B*) = P(A|B)*P(B)

Joint Probability Graphically

- P(disc golf) = 0.25 or 25%
- P(rain) = 0.20 or 20%
- P(disc golf AND rain) = 20% chance of rain, 5% of that time I play disc golf→5% x 20% = 1%

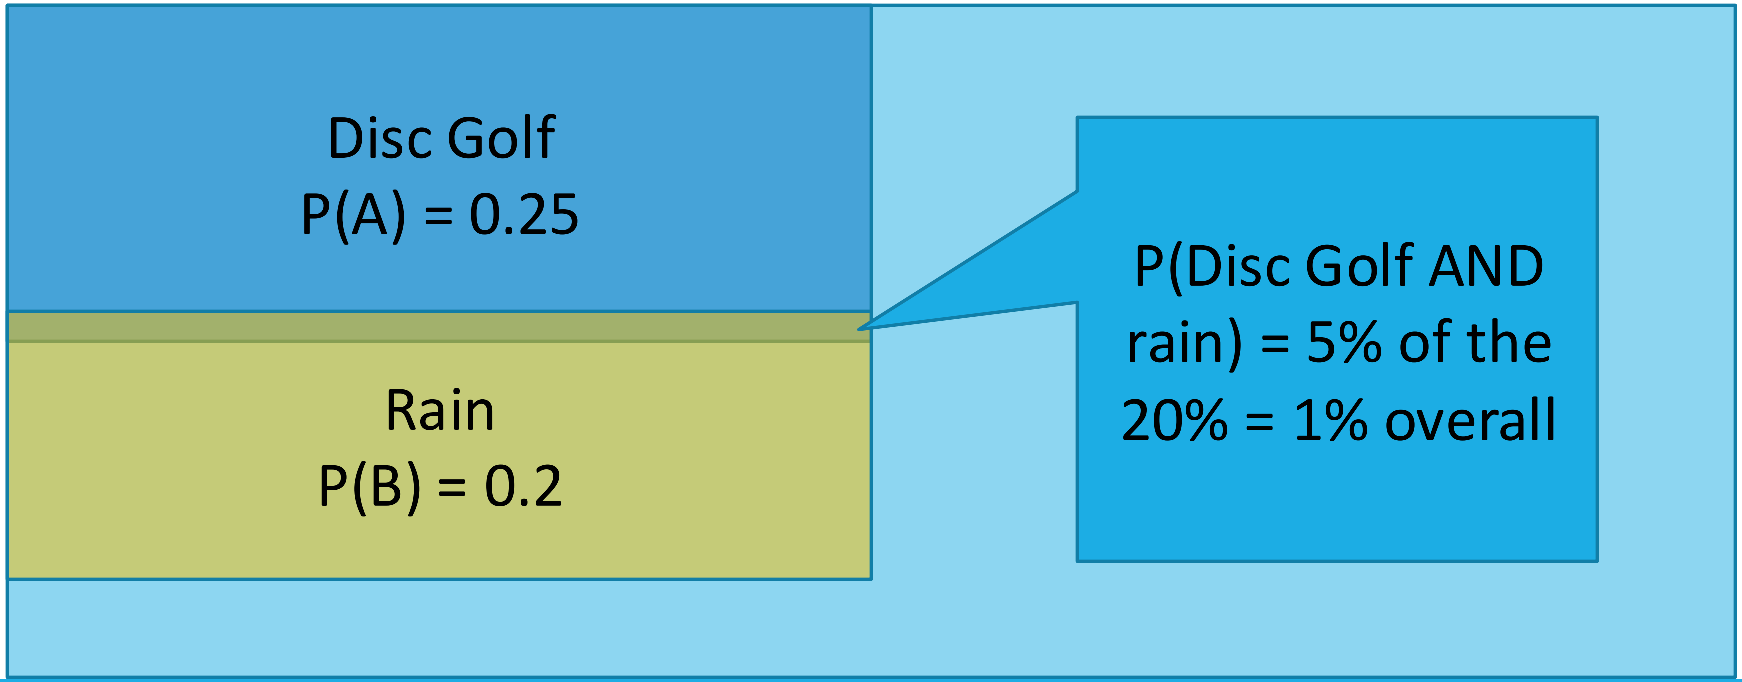

Conditional Probability Graphically

- P(Disc Golf|rain) = 0.05 or 5%
  - Focus on (condition on, “given”) orange “rain” space
  - Percent of orange space covered by blue

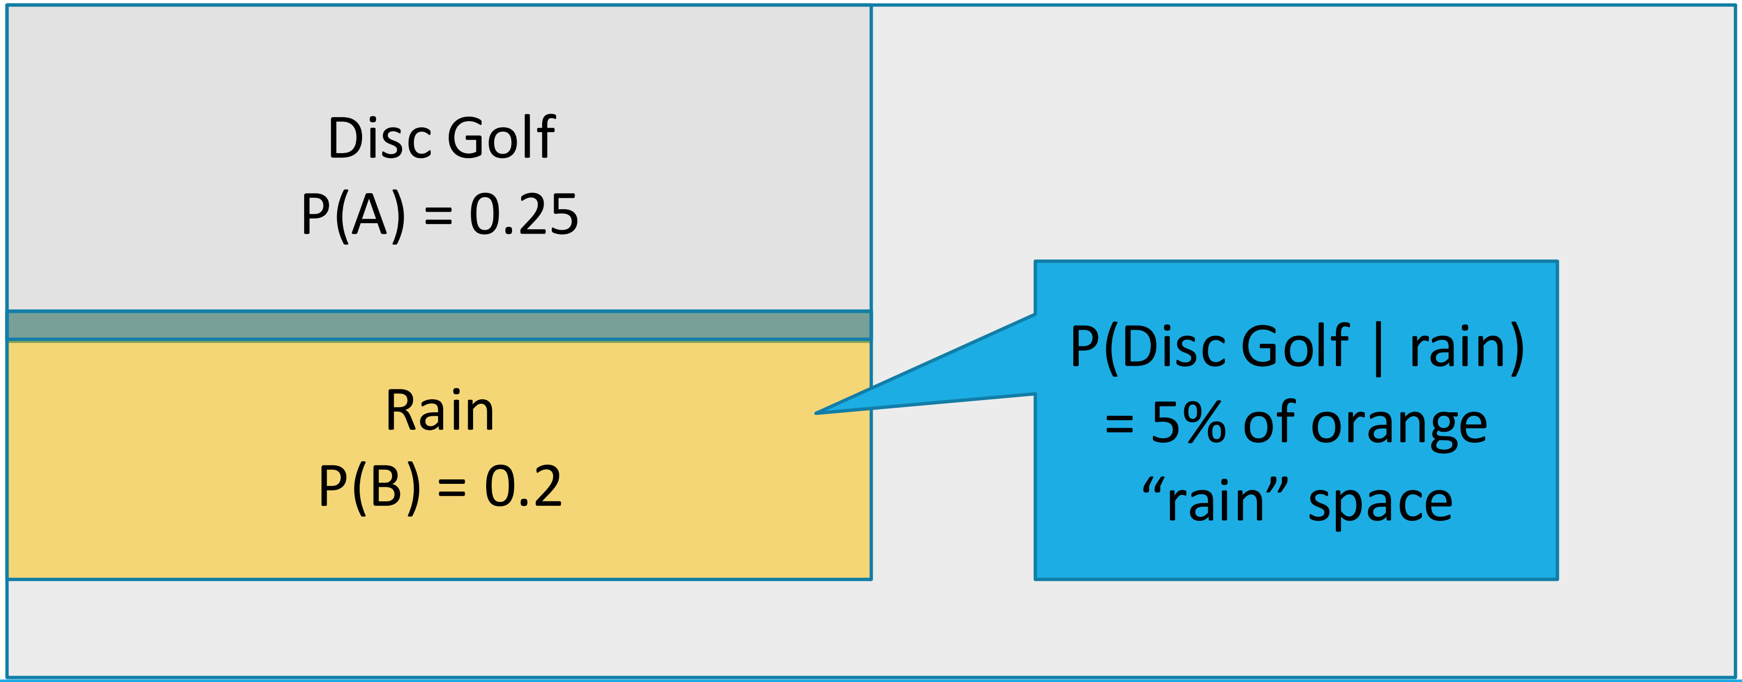

Relating Joint and Conditional Probabilities

- Relate conditional and joint probabilities:
  - **Joint Probability** = $\boldsymbol{P(A \text{ and } B) = P(A|B) * P(B)}$
  - Proportion of orange square covered by blue x size of orange square

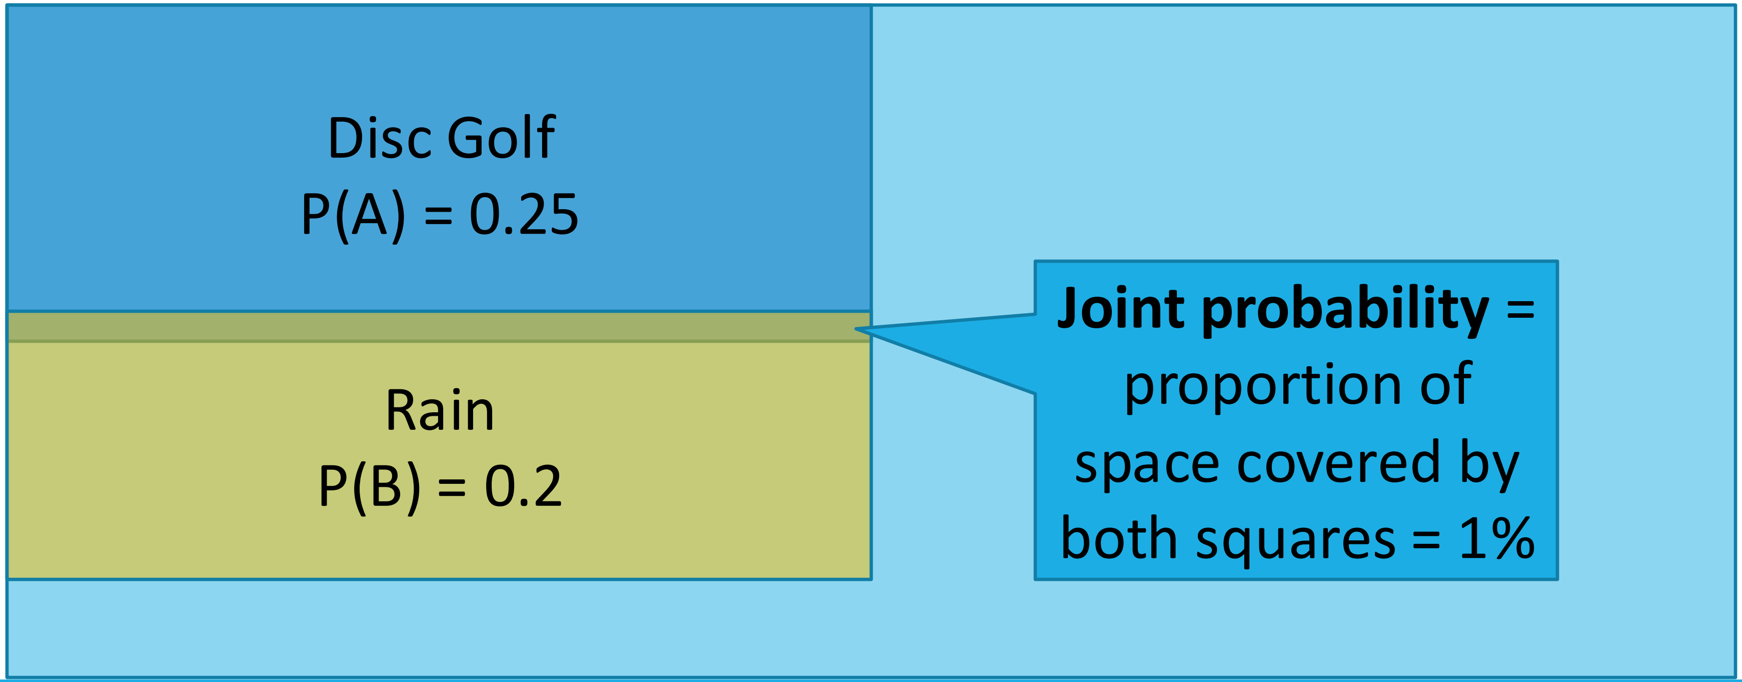

Relating Joint and Conditional Probabilities

- Relate conditional and joint probabilities:
  - **Conditional Probability =** $\boldsymbol{P(A|B) = \frac{P(A \text{ and } B)}{P(B)}}$
  - Percentage of orange square covered by blue

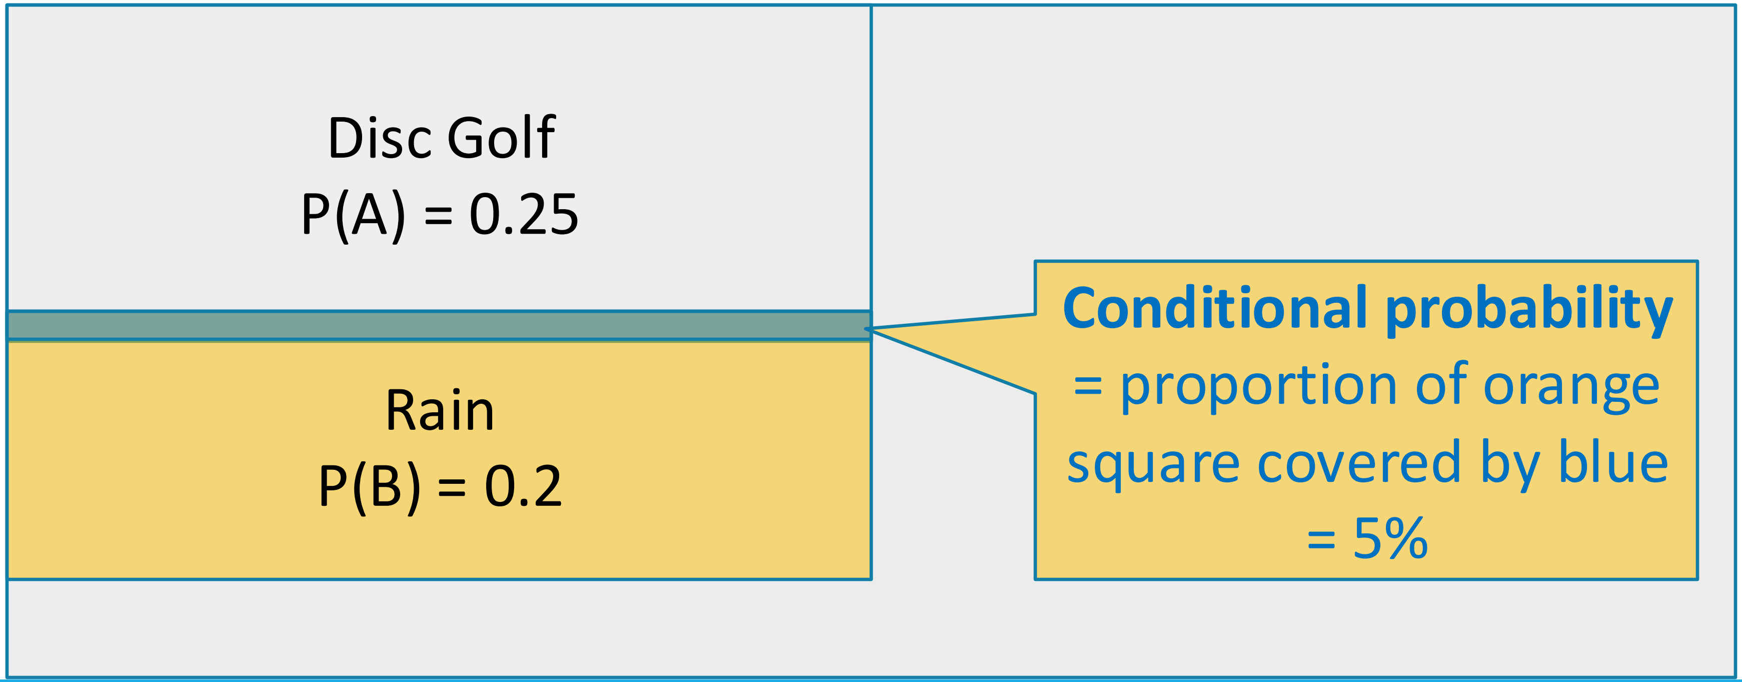

Relating Marginal and Conditional Probabilities

- **Law of Total Probability or Chain Rule:**
  - If you have an event B that takes on *n* values and an event A conditional on it,

$$\boldsymbol{P(A) = \sum_{n} P(A|B_n) * P(B_n)}$$

  - It's a simple weighted average

- What is the relationship between marginal and conditional probabilities?
  - P(Disc Golf)
    = P(Disc Golf|rain)*P(rain) + P(Disc Golf|no rain)*P(no rain)
    = 0.05*0.2 + 0.30*0.8 = 0.01 + 0.24 = 0.25

Relating Marginal and Conditional Probabilities

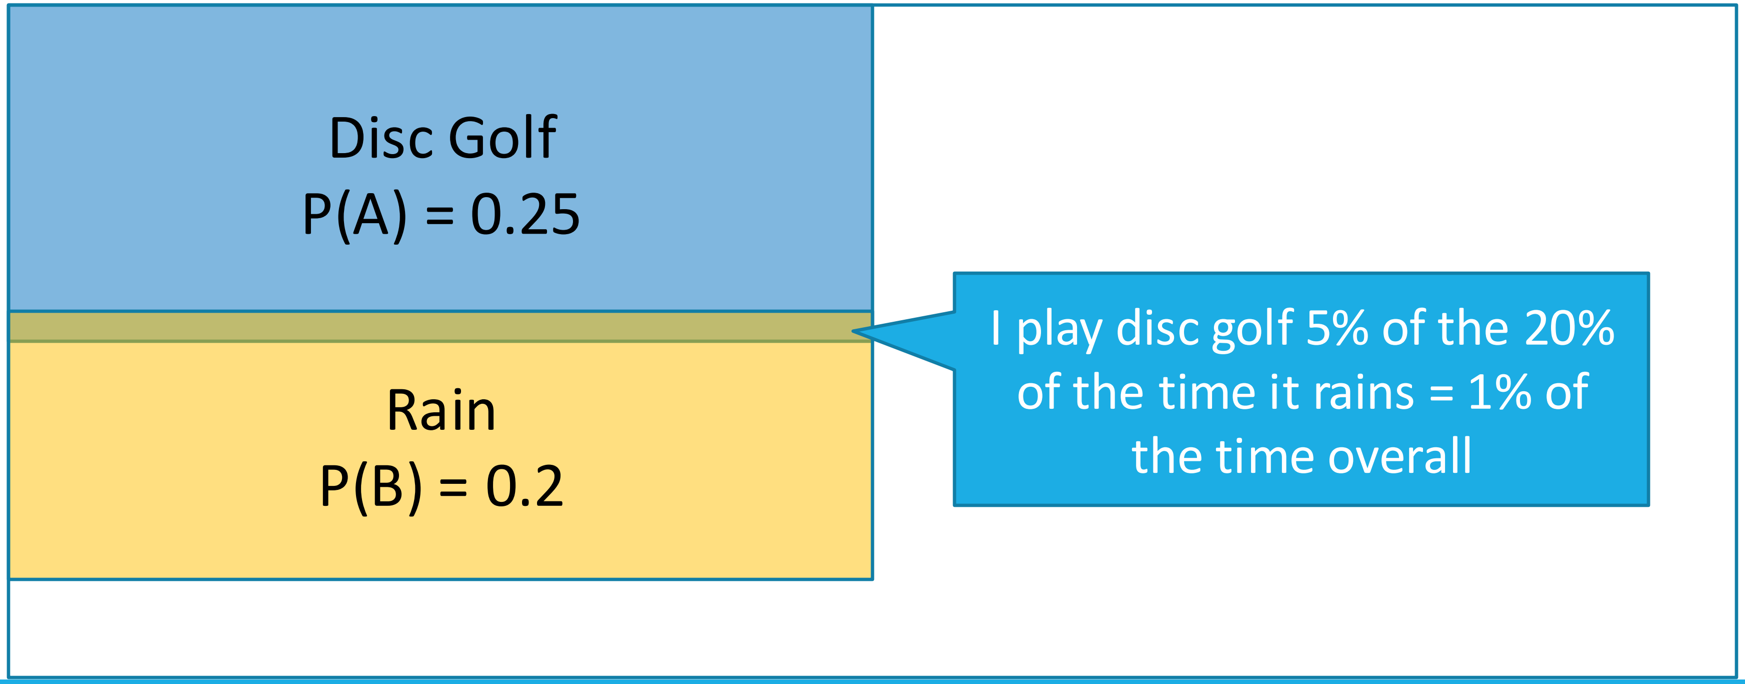

Relating Marginal and Conditional Probabilities

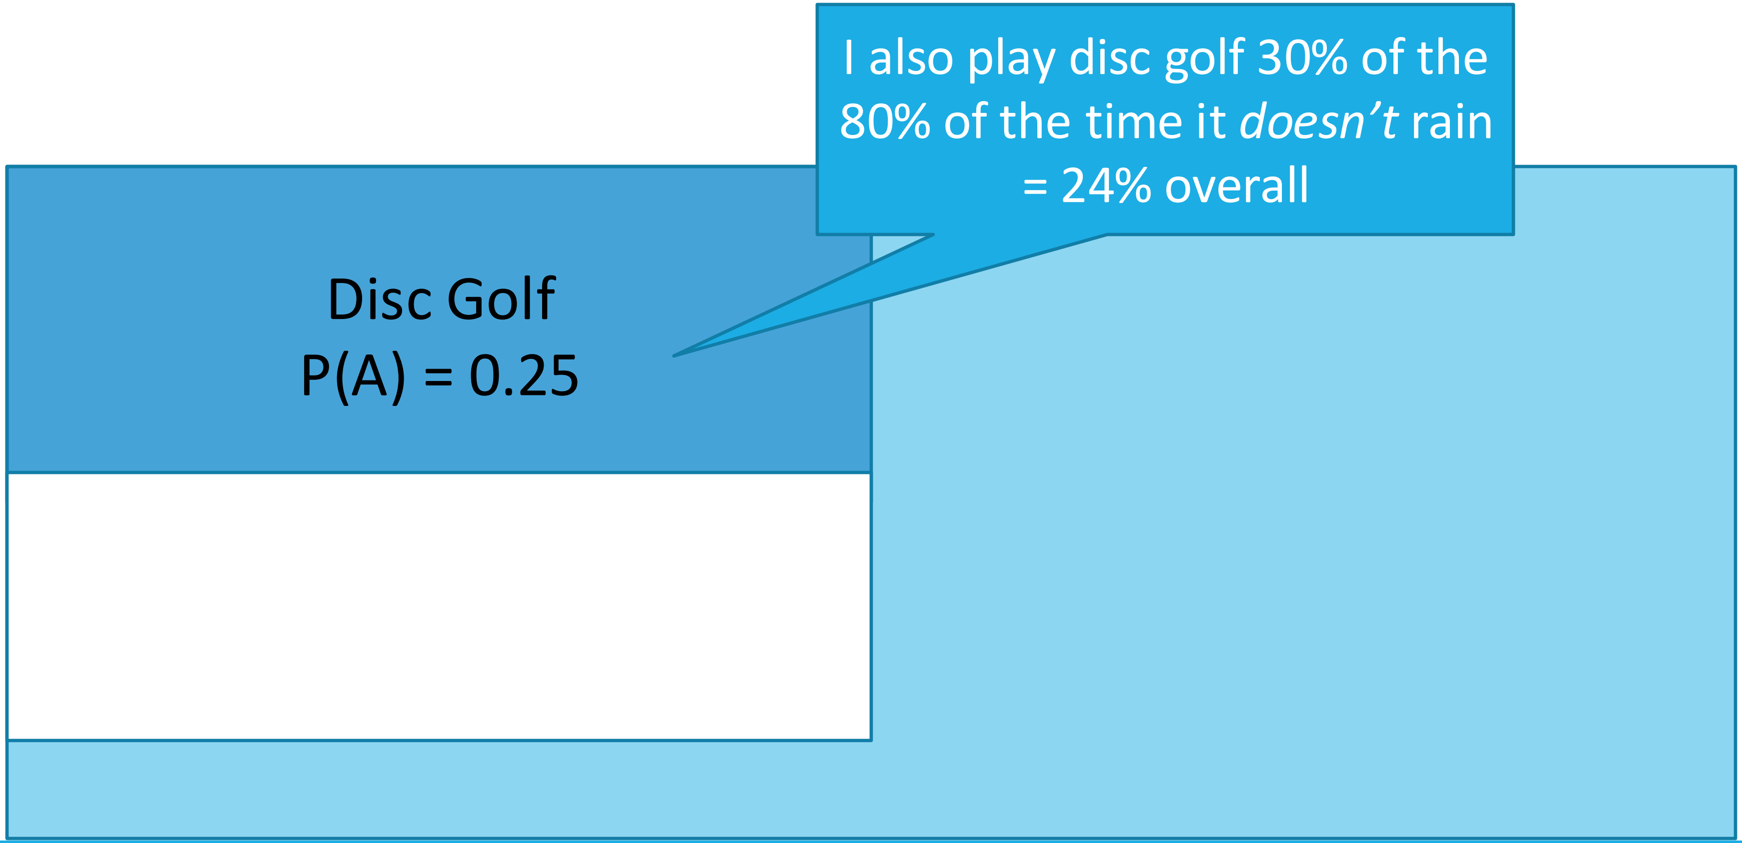

Relating Marginal and Conditional Probabilities

- **Total probability**, P(Disc Golf) = 1% + 24% = 25%

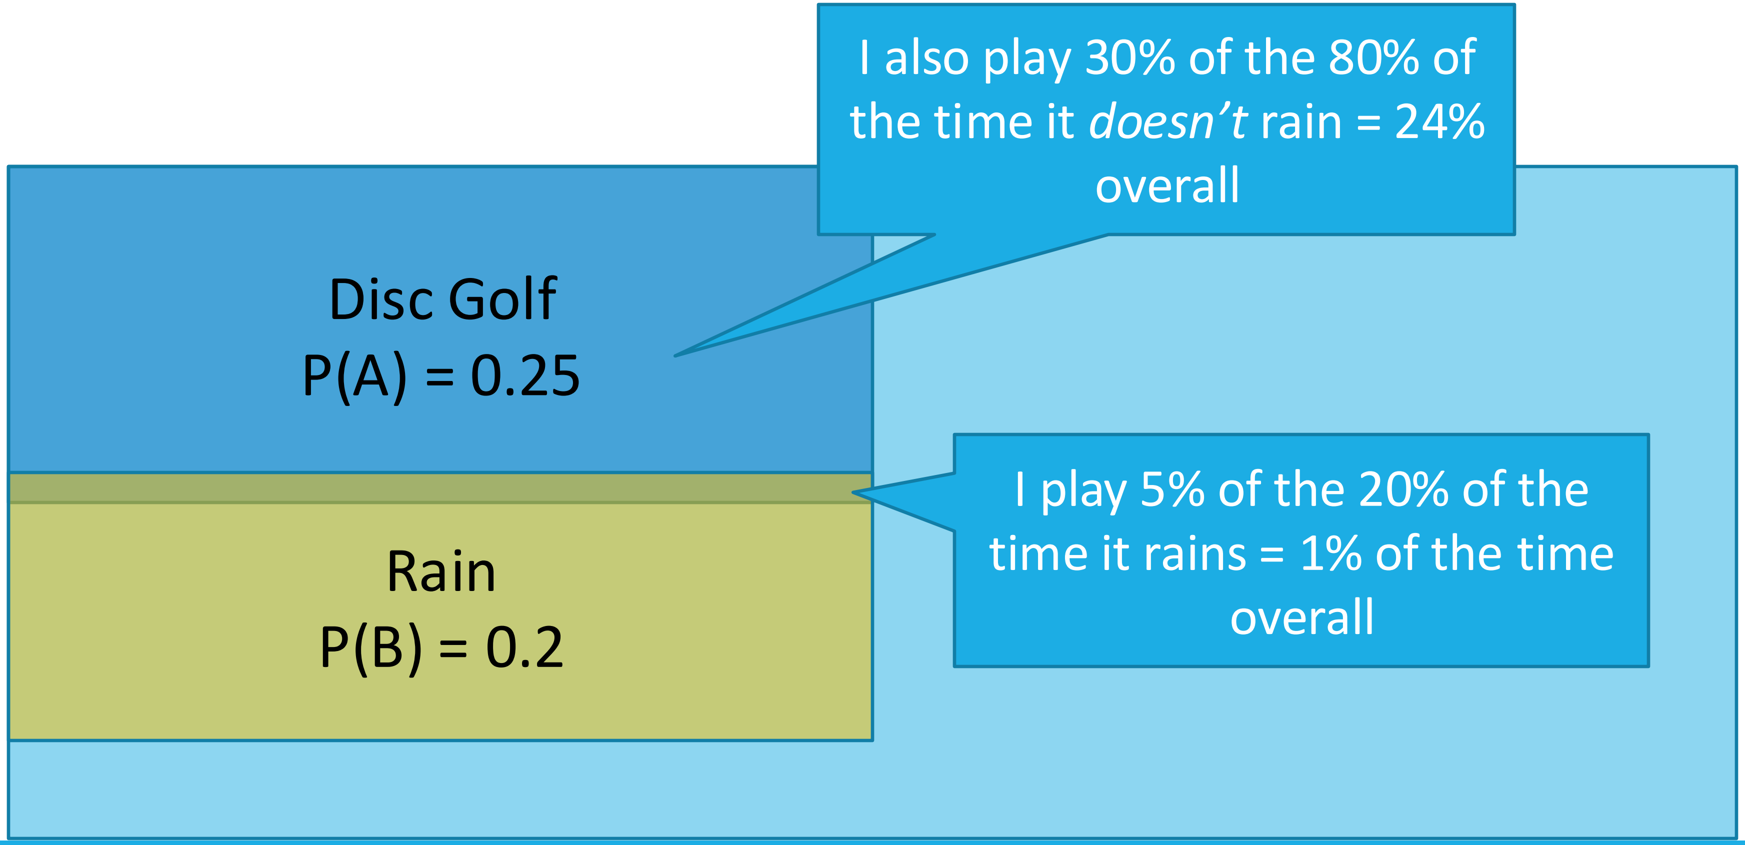

Conditional Probability

- Conditional probabilities crop up everywhere
  - Trials: What's the probability the defendant is guilty given the evidence?
    - P(Guilty | Evidence)
  - Medical tests: What's the probability you have a disease given a test is positive?
    - P(Sick | Test +)
  - And much more…

Flipping Conditional Probabilities

Let’s go back to our initial problem:

- What we want:

$$\Pr(Innocent \mid Matching)$$

- What we got:

$$\Pr(Matching \mid Innocent)$$

Using what we learned about conditional probability, can we use one to get the other?

Flipping Conditional Probabilities

Put another way:

We have $\Pr(A \mid B)$ and we want $\Pr(B \mid A)$

We know:

$$\Pr(B \mid A) = \frac{\Pr(A\ and\ B)}{\Pr(A)}$$

By the **chain rule:**

$$\Pr(B \mid A) = \frac{\Pr(A \mid B)\Pr(B)}{\Pr(A)}$$

This is called **Bayes' rule**.

Flipping Conditional Probabilities

Plugging in for our problem:

$$\Pr(Innocent \mid Matching) = \frac{\Pr(Matching \mid Innocent)\Pr(Innocent)}{\Pr(Matching)}$$

We already know that $\Pr(Matching \mid Innocent) \approx \frac{1}{500{,}000}$

Another piece of info, let's assume that the eye witness' account was perfect:

$$P(Matching \mid Guilty) = 1$$

We need one last piece of info to get rolling – the **prior** probability that the couple is innocent.

Flipping Conditional Probabilities

What is a **prior probability**?

- Before learning any new info about the couple, what is our **belief** that the couple is innocent?
- Sometimes, this probability comes from knowledge about the problem
- Other times, this probability is inherent to the problem.

What should be our prior here - $\Pr(Innocent)$?

- Let's assume that there are 2,000,000 couples in LA
- One is guilty
- So, $\Pr(Guilty) = \frac{1}{2{,}000{,}000}$ and $\Pr(Innocent) = \frac{1{,}999{,}999}{2{,}000{,}000}$

Flipping Conditional Probabilities

Plugging in for our problem:

$$\Pr(Innocent \mid Matching) = \frac{\Pr(Matching \mid Innocent)\Pr(Innocent)}{\Pr(Matching)}$$

$$\Pr(Innocent \mid Matching) = \frac{\frac{1}{500{,}000} \times \frac{1{,}999{,}999}{2{,}000{,}000} \approx .000002}{\Pr(Matching)}$$

By the law of total probability:

$$\Pr(Matching) = \Pr(Matching \mid Innocent)\Pr(Innocent) + \Pr(Matching \mid Guilty)\Pr(Guilty)$$

$$\Pr(Matching) = \frac{1}{500{,}000} \times \frac{1{,}999{,}999}{2{,}000{,}000} + 1 \times \frac{1}{2{,}000{,}000} \approx .0000025$$

Flipping Conditional Probabilities

Plugging in for our problem:

$$\Pr(Innocent \mid Matching) = \frac{\Pr(Matching \mid Innocent)\Pr(Innocent)}{\Pr(Matching)}$$

$$\Pr(Innocent \mid Matching) = \frac{.000002}{.0000025} = \frac{4}{5}$$

Using our very reasonable prior and assuming that the eyewitness account is 100% perfect, there is an 80% chance that the couple is innocent!

Flipping Conditional Probabilities

Our analysis said that there's an 80% chance that the couple is innocent.

The analyst said that the probability that couple was innocent was essentially 0!

This discrepancy arises because the mathematician asked the wrong question!

- It's very unlikely that a randomly selected innocent person would match the description of the criminals!

But that's the wrong question!

- Having found a couple that does match the description of the criminals, how likely is it that they are innocent?
- Because there are soooo many couples in LA, there are likely 5 that match that description!
- If that's the only info, then we can't claim guilt – it's more likely than not that they aren't guilty!

Flipping Conditional Probabilities

But, Dr. McAlister – I think your prior probability of guilt is wrong. You're cooking the game in the favor of the defendants.

Let's think about how our results might change if our **prior** changes.

Things we'll take at face value:

- $\Pr(Matching \mid Innocent) = \frac{1}{500{,}000}$
- $\Pr(Matching \mid Guilty) = 1$

Let x be the prior probability that the couple is innocent, then:

$$\Pr(Innocent \mid Matching) = \frac{\Pr(Matching \mid Innocent) \times x}{\Pr(Matching \mid Innocent) \times x + \Pr(Matching \mid Guilty) \times (1 - x)}$$

Flipping Conditional Probabilities

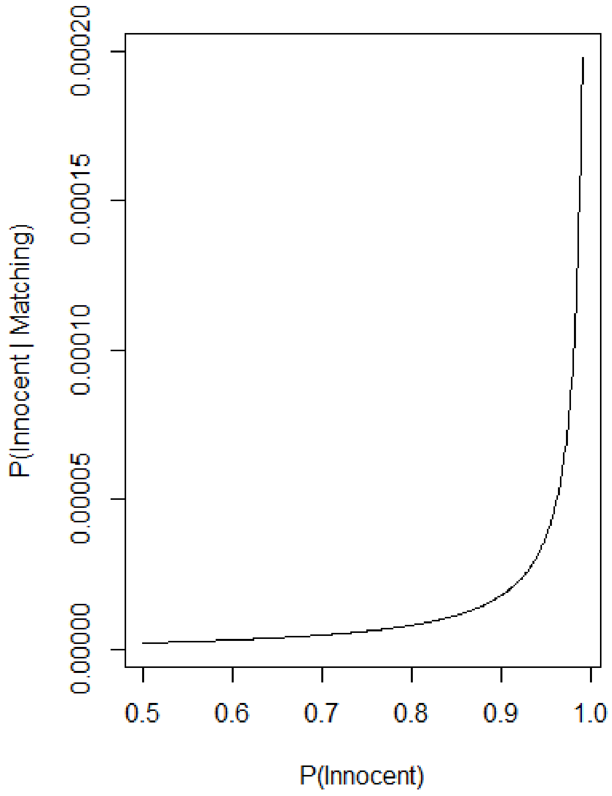

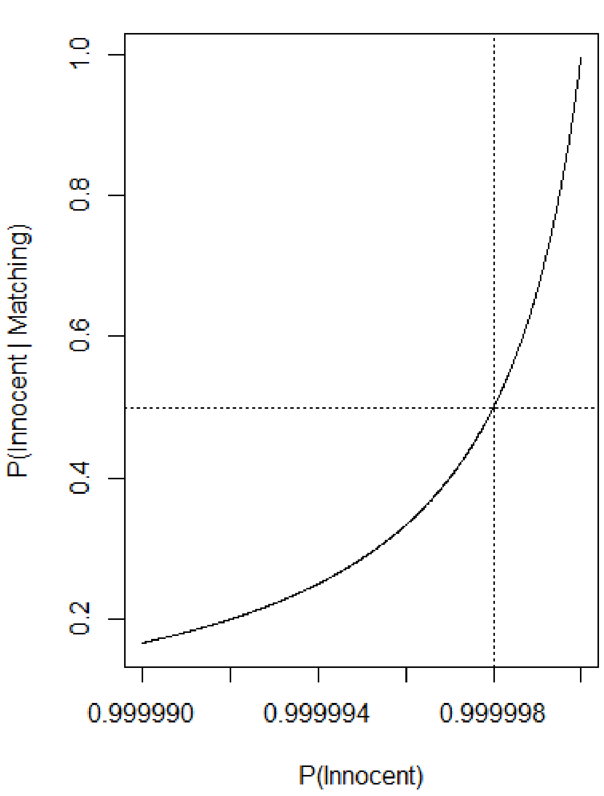

Flipping Conditional Probabilities

Right around
1 in 500,000
to be 50/50

The burden to show a high likelihood of guilt is low in this case – we just need to be able to say that our prior belief of innocence is less than .999998.

- But, that's hard to hit without more information
- Working purely on matching physical characteristics is not sufficient!

$P(A|B) \neq P(B|A)$ in most circumstances!

- The info is not symmetric! We need to be really careful about equating the two.
- The difference in these quantities can be the difference between freedom and life behind bars!
- It is **critical** to answer the question you actually want to answer and not conflate the two probabilities!

Positive Test Results

Consider the following test for a disease:

False Positive Rate of 5%

- If you don’t have the disease the test correctly returns a negative result 99% of the time
  - P(-| Don’t have Disease) = 95%
  - P(+| Don’t have Disease) = 5% (FALSE POSITIVES)
- In statistics lingo Type I error

False Negative Rate of 10%

- If you do have the disease the test correctly returns a positive result 95% of the time
  - P(+|Disease) = 90%
  - P(-|Disease) = 10% (FALSE NEGATIVES)
- In statistics lingo: Type II error

Positive Test Results

Let’s consider two **prevalences/base rates** among the population:

- For the first, 1% (.01 probability) of people actually have the disease (Celiac's disease)
- For the second, 20% (.2 probability) of people actually have the disease (COVID given cough and fever)

Uh oh. You've gotten a positive on an at-home COVID test when you have a cough and a fever! Do you actually have COVID?

Question:

$$P(\text{COVID} \mid +)$$

Positive Test Results

$$P(\text{COVID}|+) = \frac{P(+|\text{COVID})P(\text{COVID})}{P(+)}$$

$$P(\text{COVID}|+) = \frac{P(+|\text{COVID})P(\text{COVID})}{P(+\text{ and COVID}) + P(+\text{ and Not COVID})}$$

$$P(\text{COVID}|+) = \frac{P(+|\text{COVID})P(\text{COVID})}{P(+|\text{COVID})P(\text{COVID}) + P(+|\text{Not COVID})P(\text{Not COVID})}$$

Positive Test Results

$$P(\text{COVID}|+) = \frac{P(+|\text{COVID})P(\text{COVID})}{P(+|\text{COVID})P(\text{COVID}) + P(+|\text{Not COVID})P(\text{Not COVID})}$$

P(COVID)	P(Not COVID)
.2	.8

P(+|COVID)	P(-|COVID)
.9	.1
P(+|Not COVID)	**P(-|Not COVID)**
.05	.95

$$P(\text{COVID}|+) = \frac{.9 \times .2}{.9 \times .2 + .05 \times .8} \approx .82$$

Positive Test Results

Your turn!

What is the probability that you have Celiac's if your at-home test gives a positive?

P(Celiac)	P(Not Celiac)
.01	.99

P(+|Celiac)	P(-|Celiac)
.9	.1
P(+|Not Celiac)	**P(-|Not Celiac)**
.05	.95

$$P(\text{Celiac's}|+) = \frac{P(+|\text{Celiac's})P(\text{Celiac's})}{P(+)}$$

Positive Test Results

$$P(+) = P(+ \text{ and Celiac's}) + P(+ \text{ and Not Celiac's})$$

$$P(+) = P(+|\text{Celiac's})P(\text{Celiac's}) + P(+|\text{Not Celiac's})P(\text{Not Celiac's})$$

Positive Test Results

$$P(+|\text{Celiac's})P(\text{Celiac's}) = .9 \times .01 = .009$$

$$P(+|\text{Not Celiac's})P(\text{Not Celiac's}) = .05 \times .99 = .0495$$

$$P(\text{Celiac's}|+) = \frac{.009}{.009 + .0495} \approx .15$$

Positive Test Results

Let’s consider two **prevalences/base rates** among the population:

- For the first, 1% (.01 probability) of people actually have the disease (Celiac's disease)
- For the second, 20% (.2 probability) of people actually have the disease (COVID given cough)

Probability your positive means you have the disease:

- COVID = .82
- Celiac's = .15

What's the purpose of an at-home Celiac's test then?!?!

- Why are these two numbers so different?

Our Gut Instinct

80% of librarians are shy. Only 40% of people who work in retail are shy. 1/10000 people are librarians. 1/100 work in retail. Steve is shy. Is Steve more likely to be a librarian or work in retail?

Base Rate Fallacy

The interpretation of the COVID numbers in the previous example is called the **base rate fallacy**:

- People interpret evidence at face value without thinking about the base information they've been given.

Example: 80% of librarians are shy. Only 40% of people who work in retail are shy. 1/10000 people are librarians. 1/100 work in retail. Steve is shy. Is Steve more likely to be a librarian or work in retail?

- Gut reaction – Steve is shy and librarians are shy. Therefore, Steve must be a librarian.
- Correct answer:

$$\Pr(Librarian \mid Shy) = \frac{P(Shy \mid Librarian)P(Librarian)}{P(Shy)} \quad ; \quad P(Retail \mid Shy) = \frac{P(Shy \mid Retail)P(Retail)}{P(Shy)}$$

$$P(Librarian \mid Shy) = \frac{.8 \times .0001}{P(Shy)} \, ; P(Retail \mid Shy) = \frac{.4 \times .01}{P(Shy)}$$

$.00008 < .004$ so we should guess Retail

Base Rate Fallacy

Examples you can think of?

Base Rate Fallacy

Examples you can think of?

Lots of examples out in the wild:

- COVID vaccine effectiveness
- Selling off stock after a market crash
- Gambler's fallacy
- At-home disease tests

When asking questions, make sure to consider the prior information you know:

- Rare events are rare! The burden of evidence is higher when the base rate is incredibly low!
- Without strong probability altering info, common events are expected. Don't be surprised when something happens because it's frequent – even if you thought it would!

Takeaways

Probability just encodes the likelihood (or relative frequency) of occurrence of events

We can talk about conditional probabilities (how likely is A to happen if I observe that B happened) or marginal probabilities (disregarding B, how likely is A to happen).

The rules of probability can be used to derive **Bayes' Theorem:**

$$\Pr(B \mid A) = \frac{P(A \mid B)P(B)}{P(A)}$$

Takeaways

Bayes' theorem can be used to translate between conditional probabilities

- But, we must know the **base rate** before making this conversion
- In general, the two conditionals are not equivalent! In fact, they can be very far off from one another!

This conflation often leads to the **base rate fallacy**:

- People interpret one conditional as being evidence for the other
- Forget about the base rates and overinterpret new information
- Can lead to incorrect conclusions!

Flipping Conditional Probabilities

Let’s work through another example that is very relevant to today’s world.

Suppose we’re interested of figuring out that efficacy of the COVID vaccine in the real world

- RCTs show that the vaccine is 90%-95% effective at preventing serious side effects of COVID that require hospitalization
- Also, pretty high numbers for preventing infection
- Is this the same out in the real world?
- Reasons to doubt: Complier average treatment effects, lack of external validity

Flipping Conditional Probabilities

Evidence: 67% of severe COVID infections in Iceland were in people who were fully vaccinated!

- This sounds really bad for the COVID vaccine – why get vaccinated if 67% of current COVID cases in the hospital are in those who are already fully vaccinated?

Why is this piece of information potentially misleading?

Flipping Conditional Probabilities

Evidence: 67% of severe COVID infections in Iceland were in people who were fully vaccinated!

- This sounds really bad for the COVID vaccine – why get vaccinated if 67% of current COVID cases in the hospital are in those who are already fully vaccinated?

What this number gives us:

$$\Pr(Vaccinated \mid Hospital)$$

What we want:

$$\Pr(Hospital \mid Vaccinated)$$

Flipping Conditional Probabilities

Let's say that our population is 71% vaccinated, like the estimate in Iceland. 366,000 people.

- 260 thousand vaccinated
- 106 thousand unvaccinated

Let's say that there are 1000 people in the hospital:

- 670 vaccinated
- 330 unvaccinated

Proportions:

- 670/260000 = .0025
- 330/106000 = .0031

In this example, the unvaccinated are 1.25x more likely to go to the hospital than the vaccinated!

Flipping Conditional Probabilities

Let's formalize the problem more:

$$\Pr(Hospital \mid Vaccine) = \frac{\Pr(Vaccine \mid Hospital)P(Hospital)}{\Pr(Vaccine \mid Hospital)\Pr(Hospital) + \Pr(Vaccine \mid No\ Hospital)\Pr(No\ Hospital)}$$

To work this out, we need to know:

- $\Pr(Hospital)$ – the marginal probability that someone, regardless of vaccination status goes to the hospital for severe COVID
- $\Pr(Vaccine \mid No\ Hospital)$ – the probability that someone has the vaccine if they don't go to the hospital for severe COVID

Flipping Conditional Probabilities

Plugging in what we know:

$$\Pr(Hospital \mid Vaccine) = \frac{.67 \times P(Hospital)}{.67 \times \Pr(Hospital) + \Pr(Vaccine \mid No\ Hospital)\Pr(No\ Hospital)}$$

Let’s assume that 71% of the unhospitalized population is vaccinated (the vaccinated rate).

Flipping Conditional Probabilities

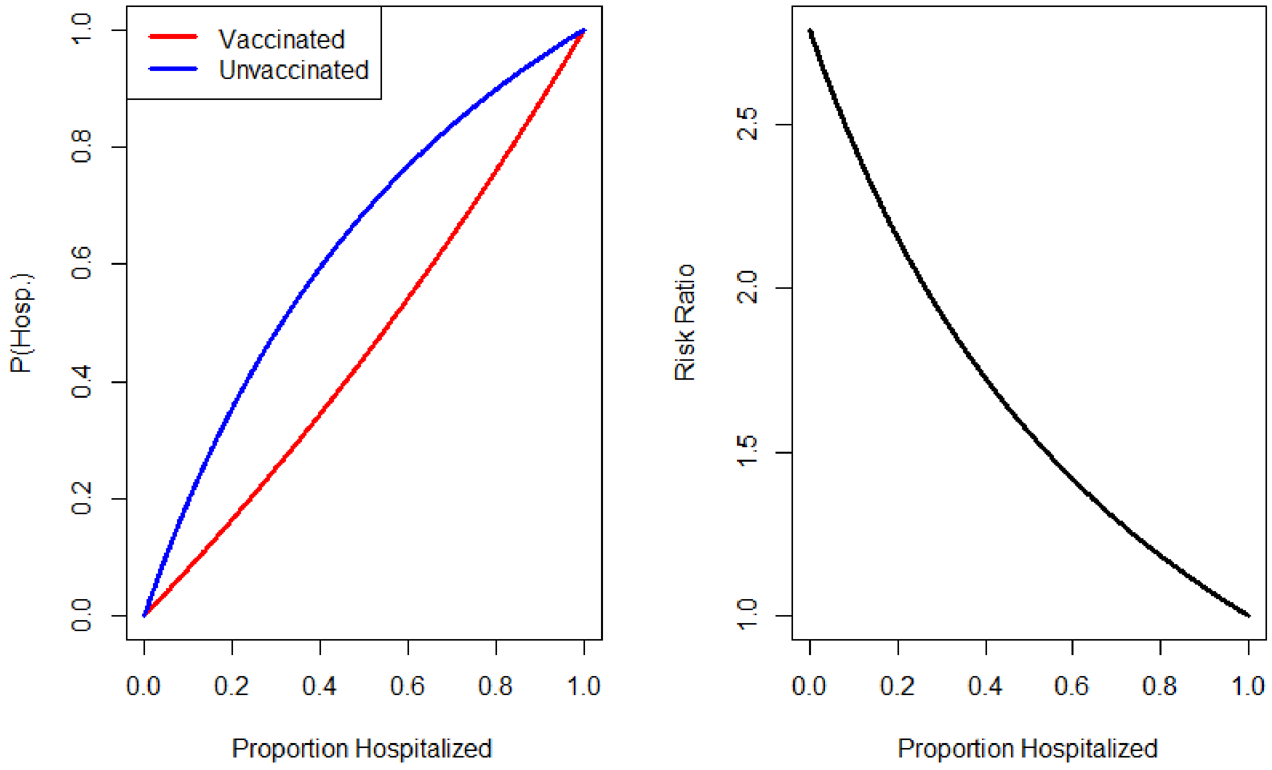

Flipping Conditional Probabilities

Depending on the proportion of individuals in the hospital, the risk ratio of being vaccinated vs. unvaccinated changes drastically!

- In world where few people are hospitalized (hospitalization is relatively rare), being vaccinated reduces the relative probability of being hospitalized **greatly**
- In a world where a lot of people are hospitalized, this is not true.

Tricky math, but easy logic:

- If everyone is getting really sick, then it's clear that the vaccine doesn't work!
- Otherwise, the difference in **base rates** means that effectiveness cannot be judged by the raw numbers, alone. Must consider how many people are vaccinated!