CSSS/SOC/STAT 321
Case-Based Statistics I

Introduction to Probability

Christopher Adolph

Department of Political Science

and

Center for Statistics and the Social Sciences

University of Washington, Seattle

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events
$\mathrm{P}(A)$	the probability that event A happens out of a continuum of possible events

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events
$\mathrm{P}(A)$	the probability that event A happens out of a continuum of possible events
$\Pr(\text{not } A)$	the probability something other than A happens

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events
$\mathrm{P}(A)$	the probability that event A happens out of a continuum of possible events
$\Pr(\text{not } A)$	the probability something other than A happens
$\Pr(A \cup B)$	the probability either A or B (or both) happen ($\cup$ stands for union)

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events
$\mathrm{P}(A)$	the probability that event A happens out of a continuum of possible events
$\Pr(\text{not } A)$	the probability something other than A happens
$\Pr(A \cup B)$	the probability either A or B (or both) happen ($\cup$ stands for union)
$\Pr(A \cap B)$	the probability A and B happen *at the same time* ($\cap$ stands for intersection)

Aside on mathematical notation

$\Pr(A)$	the probability that event A happens out of a discrete set of possible events
$\mathrm{P}(A)$	the probability that event A happens out of a continuum of possible events
$\Pr(\text{not } A)$	the probability something other than A happens
$\Pr(A \cup B)$	the probability either A or B (or both) happen ($\cup$ stands for union)
$\Pr(A \cap B)$	the probability A and B happen *at the same time* ($\cap$ stands for intersection)
$\Pr(A\|B)$	the probability A happens given that B is certain to happen

Four examples

Time for a second opinion? During a routine checkup, your doctor tells you some bad news: you tested positive for a rare disease. The disease affects 1 in 10,000 people, and the test has a 99% effectiveness rate. What are the chances you have the disease?

Concepts applied: Sample spaces. Conditional probability.

Four examples

Time for a second opinion? During a routine checkup, your doctor tells you some bad news: you tested positive for a rare disease. The disease affects 1 in 10,000 people, and the test has a 99% effectiveness rate. What are the chances you have the disease?

Concepts applied: Sample spaces. Conditional probability.

A statistician plays the lottery What are your chances of winning the lottery?

Concepts applied: Complex events. Independence. Joint probability. Expected value.

Four examples

The prosecutor's fallacy Out of a city of 6 million people, a prosecutor has matched your DNA to a crime scene, and tells the jury that only 1 in 1,000,000 people have matching DNA. You know you are innocent. But how can you convince a jury?

Concepts applied: Inverse probability, Bayes Rule.

Four examples

The prosecutor's fallacy Out of a city of 6 million people, a prosecutor has matched your DNA to a crime scene, and tells the jury that only 1 in 1,000,000 people have matching DNA. You know you are innocent. But how can you convince a jury?

Concepts applied: Inverse probability, Bayes Rule.

Let's make a deal On a game show, you get to choose one of three doors. Only one has a car; the others have goats. After you suggest a door, the host opens one door (to reveal a goat), and offers you the chance to switch to door number 3. Do you stay or switch?

Concepts applied: Monte Carlo simulation for complex probability problems.

Essential concepts

Event Any specific outcome that might occur. Mutually exclusive with other events.

Example: A World Cup Soccer game between the US and England could end in a victory for the US.

Essential concepts

Event Any specific outcome that might occur. Mutually exclusive with other events.

Example: A World Cup Soccer game between the US and England could end in a victory for the US.

Sample space A set of *all* the events that might occur.

Example: A World Cup Soccer game between the US and England could end in three ways: Win, Loss, Tie

Essential concepts

Event Any specific outcome that might occur. Mutually exclusive with other events.

Example: A World Cup Soccer game between the US and England could end in a victory for the US.

Sample space A set of *all* the events that might occur.

Example: A World Cup Soccer game between the US and England could end in three ways: Win, Loss, Tie

Probability A number, between 0 and 1, assigned to each of the possible outcomes in the sample space.

Example: $\mathrm{Pr}(\mathrm{Win}) = 0.3$, $\mathrm{Pr}(\mathrm{Loss}) = 0.5$, $\mathrm{Pr}(\mathrm{Tie}) = 0.2$

Axioms of Probability

An axiom is an assumption we cannot prove, and must make to get started in a field of mathematics.

Probability theory relies on just three axioms:

1. Each event has a probability of 0 or more. (Anything could happen.)

Axioms of Probability

An axiom is an assumption we cannot prove, and must make to get started in a field of mathematics.

Probability theory relies on just three axioms:

1. Each event has a probability of 0 or more. (Anything could happen.)
2. The total probability of all the events in the sample space is 1. (Something must happen.)

Axioms of Probability

An axiom is an assumption we cannot prove, and must make to get started in a field of mathematics.

Probability theory relies on just three axioms:

1. Each event has a probability of 0 or more. (Anything could happen.)
2. The total probability of all the events in the sample space is 1. (Something must happen.)
3. If two events A and B cannot simultaneously occur, the probability that either occurs is $\Pr(A) + \Pr(B)$

Approaches to Probability

Frequency Interpretation We observe how often an event occurs in a set of trials. The ratio of successes to trials show reflect the long-term probability of the event.

Approaches to Probability

Frequency Interpretation We observe how often an event occurs in a set of trials. The ratio of successes to trials show reflect the long-term probability of the event.

Theoretical interpretation The above suggests there is also some true, usually unknown, probability that an event occurs in nature.

Approaches to Probability

Frequency Interpretation We observe how often an event occurs in a set of trials. The ratio of successes to trials show reflect the long-term probability of the event.

Theoretical interpretation The above suggests there is also some true, usually unknown, probability that an event occurs in nature.

Subjective Interpretation Probabilities may also reflect personal beliefs about the likelihood of an event.

We'll mostly rely on the frequency interpretation, but will use all three at different points in the course

Visualizing sample spaces and events with Venn Diagrams

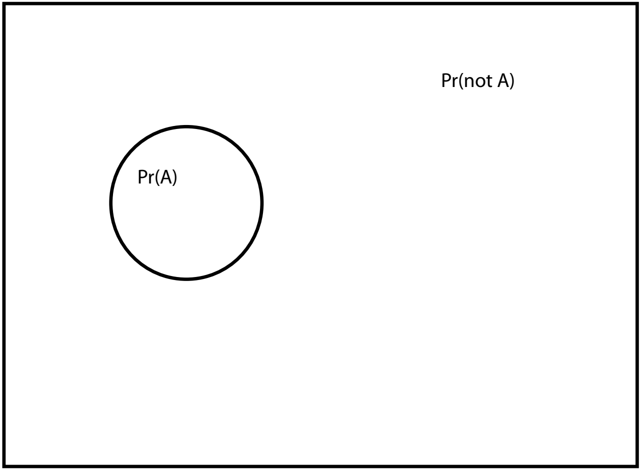

The entire rectangle is the *sample space*, and has area, or total probability, of 1

Visualizing sample spaces and events with Venn Diagrams

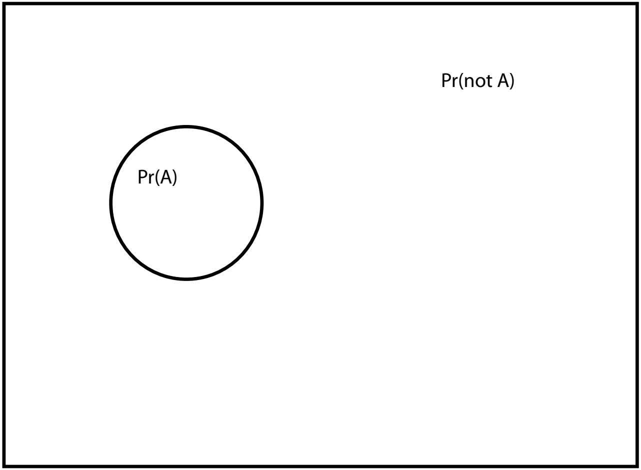

The entire rectangle is the *sample space*, and has area, or total probability, of 1

The area of the small circle represents the probability that event A happens

Visualizing sample spaces and events with Venn Diagrams

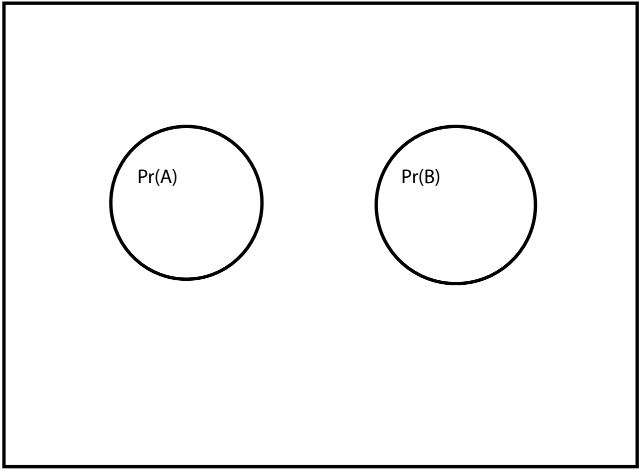

If the circles for two different events don't overlap, they are disjoint

Visualizing sample spaces and events with Venn Diagrams

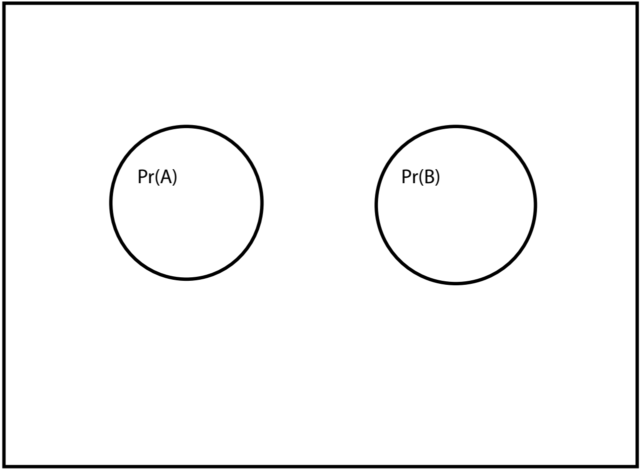

If the circles for two different events don't overlap, they are disjoint

We say these events are mutually exclusive

Visualizing sample spaces and events with Venn Diagrams

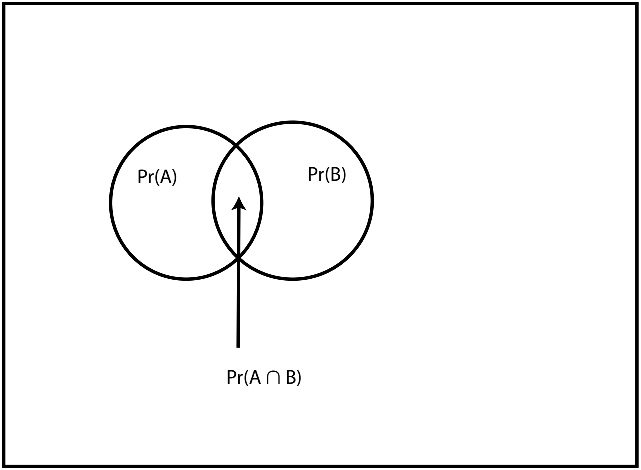

If the circles overlap, both events could happen at once

Visualizing sample spaces and events with Venn Diagrams

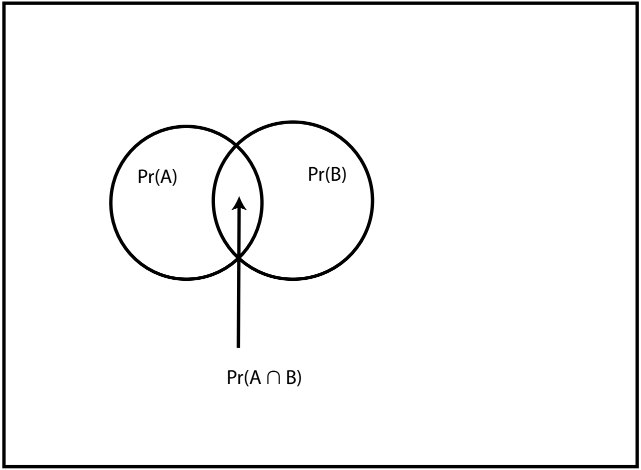

If the circles overlap, both events could happen at once

The probability both happen is the *joint* probability of $A \cap B$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not}\, A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Addition rule (mutually exclusive events) $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $\Pr(A \cap B) = 0$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Addition rule (mutually exclusive events) $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $\Pr(A \cap B) = 0$

Conditional probability $\Pr(A|B) = \Pr(A \cap B)/\Pr(B)$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Addition rule (mutually exclusive events) $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $\Pr(A \cap B) = 0$

Conditional probability $\Pr(A|B) = \Pr(A \cap B)/\Pr(B)$

Independence of events Events are independent *only if* $\Pr(A|B) = \Pr(A)$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Addition rule (mutually exclusive events) $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $\Pr(A \cap B) = 0$

Conditional probability $\Pr(A|B) = \Pr(A \cap B)/\Pr(B)$

Independence of events Events are independent *only if* $\Pr(A|B) = \Pr(A)$

Joint probability (independent events) $\Pr(A \cap B) = \Pr(A)\Pr(B)$

Tools to solve probability problems

Complement rule $\Pr(A) = 1 - \Pr(\text{not } A)$

Addition rule (generally) $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$

Addition rule (mutually exclusive events) $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $\Pr(A \cap B) = 0$

Conditional probability $\Pr(A|B) = \Pr(A \cap B)/\Pr(B)$

Independence of events Events are independent *only if* $\Pr(A|B) = \Pr(A)$

Joint probability (independent events) $\Pr(A \cap B) = \Pr(A)\Pr(B)$

Joint probability (generally) $\Pr(A \cap B) = \Pr(A)\Pr(B|A)$

Time for a Second Opinion?

During a routine checkup, your doctor tells you some bad news:
you tested positive for a rare disease.

The disease affects 1 in 10,000 people,
and the test has a 99% effectiveness rate.

What are the chances you have the disease?

Steps to solve a probability problem

1. Identify the possible events
2. Identify the quantity of interest in terms of probability of specific events
3. Collect all the probabilities you know
4. Use the rules of probability to calculate what you want to know from what you do know

Identify the possible events

You can either have the disease or not have the disease

You can either test positive or negative

This leads to four possible combinations which comprise the whole sample space:

- have disease and positive
- have disease and negative
- no disease and positive
- no disease and negative

Venn Diagram of disease and test events

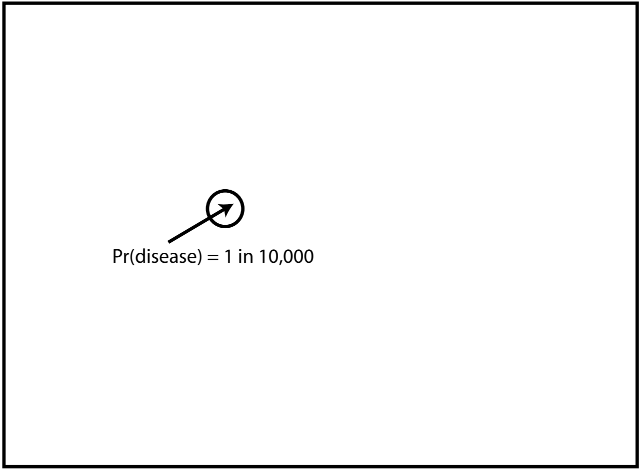

Let's start with the event of having the disease

Venn Diagram of disease and test events

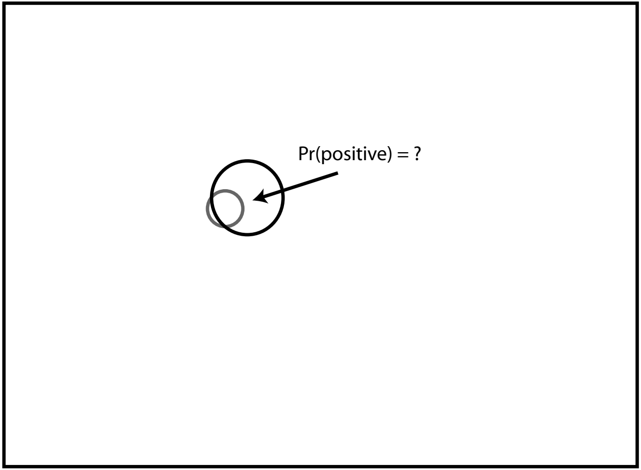

Testing positive is a separate, possibly joint event with having the disease

Venn Diagram of disease and test events

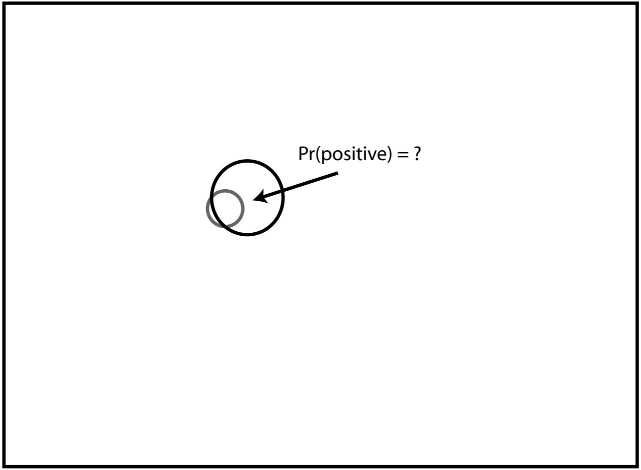

Testing positive is a separate, possibly joint event with having the disease

We don’t yet know the marginal probability of testing positive

Venn Diagram of disease and test events

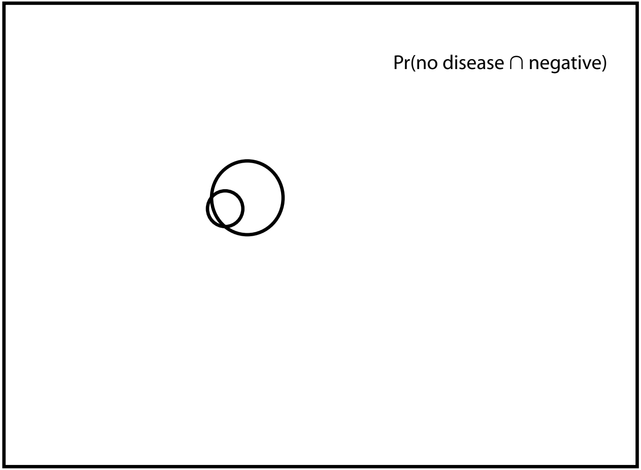

The area outside both circles is the probability of neither having the disease nor testing positive

Venn Diagram of disease and test events

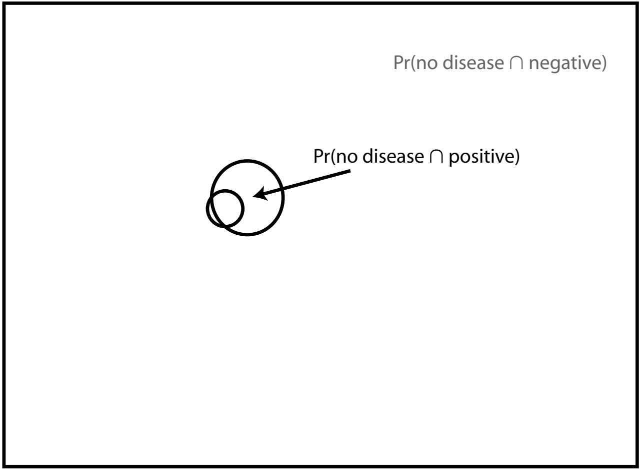

This large crescent shape is the probability of testing positive even though you don't have the disease

Venn Diagram of disease and test events

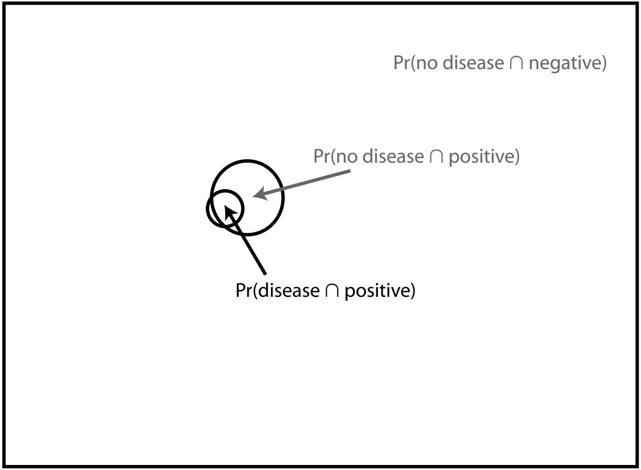

This intersection is the probability of testing positive and having the disease

Venn Diagram of disease and test events

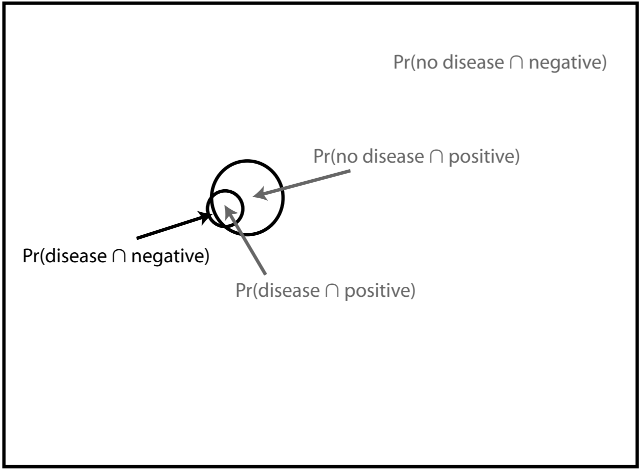

And this crescent is the probability of having the disease but failing to detect it

Identify the quantity of interest

Clearly, we want to know the possibility you have the disease *given* a positive test result

The simple probability of having a disease isn't enough – you are worried because you have new information: you had a positive test

In formal terms, we want to find:

$$\Pr(\text{disease}|\text{positive})$$

This is not the same as either $\Pr(\text{disease})$ or $\Pr(\text{positive})$

How do we find it? Let's start with what we know

Collect all the probabilities you know

We know how likely a random person is to have the disease:

$$\begin{aligned}
\Pr(\text{disease}) &= \text{1 in 10,000} \\
&= 0.01\% \\
&= 0.0001
\end{aligned}$$

Collect all the probabilities you know

We know how likely a random person is to have the disease:

$$\begin{aligned}\Pr(\text{disease}) &= \text{1 in 10,000} \\ &= 0.01\% \\ &= 0.0001\end{aligned}$$

From this, we can calculate the complement:

$$\begin{aligned}\Pr(\text{no disease}) &= 1 - \Pr(\text{disease}) \\ &= \text{9,999 in 10,000} \\ &= 99.99\% \\ &= 0.9999\end{aligned}$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is "99% effective", meaning it has a 99% probability of detecting the correct disease status, so:

$$\begin{aligned}\Pr(\text{positive}|\text{disease}) &= 0.99 \\ &= 99 \text{ in } 100\end{aligned}$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is “99% effective”, meaning it has a 99% probability of detecting the correct disease status, so:

$$\begin{aligned}
\Pr(\text{positive}|\text{disease}) &= 0.99 \\
&= 99 \text{ in } 100 \\
\Pr(\text{negative}|\text{disease}) &= 0.01 \\
&= 1 \text{ in } 100
\end{aligned}$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is “99% effective”, meaning it has a 99% probability of detecting the correct disease status, so:

$$
\begin{aligned}
\Pr(\text{positive}|\text{disease}) &= 0.99 \\
&= 99 \text{ in } 100 \\
\Pr(\text{negative}|\text{disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\Pr(\text{positive}|\text{no disease}) &= 0.01 \\
&= 1 \text{ in } 100
\end{aligned}
$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is “99% effective”, meaning it has a 99% probability of detecting the correct disease status, so:

$$\begin{aligned}
\Pr(\text{positive}|\text{disease}) &= 0.99 \\
&= 99 \text{ in } 100 \\
\Pr(\text{negative}|\text{disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\Pr(\text{positive}|\text{no disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\Pr(\text{negative}|\text{no disease}) &= 0.99 \\
&= 99 \text{ in } 100
\end{aligned}$$

What we still don't know

$$\begin{aligned}\Pr(\text{disease}|\text{positive}) &= \textbf{?}\\ \Pr(\text{no disease}|\text{positive}) &= \textbf{?}\end{aligned}$$

How are we going to calculate these?

What we still don't know

$$\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \textbf{?} \\
\Pr(\text{no disease}|\text{positive}) &= \textbf{?}
\end{aligned}$$

How are we going to calculate these?

We have a formula for conditional probabilities:

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

What we still don't know

$$\begin{aligned} \Pr(\text{disease}|\text{positive}) &= \textbf{?} \\ \Pr(\text{no disease}|\text{positive}) &= \textbf{?} \end{aligned}$$

How are we going to calculate these?

We have a formula for conditional probabilities:

$$\begin{aligned} \text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\ \Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \end{aligned}$$

So we need to calculate some joint probabilities
(the areas in the Venn Diagrams)

Use the rules of probability to calculate the missing probabilities

2×2 table for disease test example

	disease	no disease	
positive			
negative			

One way to calculate missing probabilities is to make a contingency table, and see if we can fill in the blanks

Use the rules of probability to calculate the missing probabilities

2×2 table for disease test example

	disease	no disease	
positive			
negative			
			1,000,000

Let's choose a large "sample" size, and then compute appropriate frequencies

That is, we'll work with the frequency interpretation of probability

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive			
negative			
	100	999,900	1,000,000

If the $\Pr(\text{disease})$ is 1 in 10,000,
then out of a sample of 1 million, about 100 should have the disease.

These simple probabilities fill in the *margins* of the table,
and so simple probabilities are often called *marginal probabilities*

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive	99		
negative	1		
	100	999,900	1,000,000

We know from the conditional probability of test results given disease that 99 of the 100 disease cases should be detected

This helps us fill in the joint probabilities of disease and test results

Use the rules of probability to calculate the missing probabilities

2 $\times$ 2 table for disease test example

	disease	no disease	
positive	99	9,999	
negative	1	989,901	
	100	999,900	1,000,000

Likewise, we can fill in the fraction of non-disease cases that the test should correctly identify

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive	99	9,999	10,098
negative	1	989,901	989,902
	100	999,900	1,000,000

That just leaves the marginal totals of positive and negative results, which we find by adding up the rows

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive	0.000099	0.009999	0.010098
negative	0.000001	0.989901	0.989902
	0.0001	0.99999	1.0

Dividing through by the grand sum of the table converts all the entries to probabilities

Note that unlikely last week, we want to divide by the overall sum, not the column sums. Our goal is to find the probabilities of each combination of events

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive			Pr(positive)
negative			Pr(negative)
	Pr(disease)	Pr(no disease)	Pr(any event)

Let's think about what we've done in terms of probabilities

The margins of the table have the simple probabilities of each event

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive			Pr(positive)
negative			Pr(negative)
	Pr(disease)	Pr(no disease)	Pr(any event)

And the cells of the table show the joint probability of each combination of events

Use the rules of probability to calculate the missing probabilities

2×2 table for disease test example

	disease	no disease	
positive			Pr(positive)
negative			Pr(negative)
	Pr(disease)	Pr(no disease)	Pr(any event)

We can use these cells to compute any conditional probability we want:

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

Use the rules of probability to calculate the missing probabilities

2×2 table for disease test example

	disease	no disease	
positive			Pr(positive)
negative			Pr(negative)
	Pr(disease)	Pr(no disease)	Pr(any event)

We can use these cells to compute any conditional probability we want:

$$\begin{aligned} \text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\ \Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})} \end{aligned}$$

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive	0.000099	0.009999	0.010098
negative	0.000001	0.989901	0.989902
	0.0001	0.99999	1.0

$$\Pr(\text{disease}|\text{negative}) = \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})}$$

Use the rules of probability to calculate the missing probabilities

2 $\times$ 2 table for disease test example

	disease	no disease	
positive	0.000099	0.009999	0.010098
negative	0.000001	0.989901	0.989902
	0.0001	0.99999	1.0

$$\begin{aligned} \Pr(\text{disease}|\text{negative}) &= \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})} \\ &= \frac{0.000099}{0.010098} = 0.0098 \end{aligned}$$

Use the rules of probability to calculate the missing probabilities

2 × 2 table for disease test example

	disease	no disease	
positive	0.000099	0.009999	0.010098
negative	0.000001	0.989901	0.989902
	0.0001	0.99999	1.0

If you randomly test people for the disease, and get a positive result, there are 99 chances in 100 that the person doesn't have the disease

And only a 1 in 100 chance of the person having the disease, despite the positive test!

Alternative Presentation: Probability Trees

Many students find conditional probabilities easier to understanding using tree diagrams

We will look at this problem a second way using a tree

Each node in the tree represents a random variable

Each branch represents a possible value of that variable

Tracing out each branch to the tip shows the joint probability that a set of variables come out a certain way

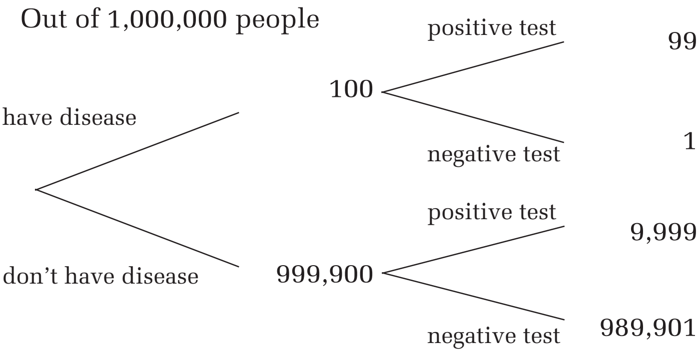

Again suppose we took 1,000,000 people, and their outcomes followed exactly the joint and marginal probabilities we determined for each event.

We would have the above tree, where numbers at the right show the total people with all the conditions on the corresponding branch

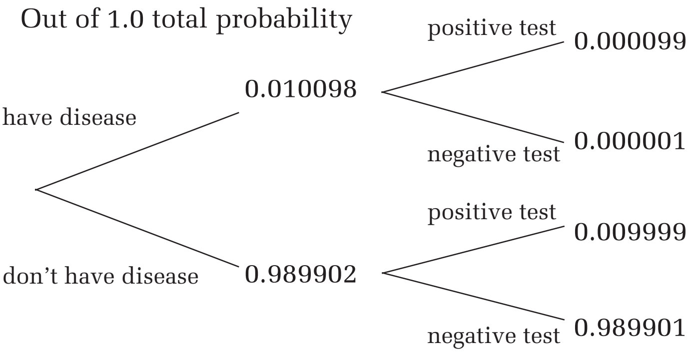

If we divide through by 1,000,000, we get the underlying probabilities.

Let’s leave them aside for the moment, and find the probability of disease given a positive diagnosis, $\Pr(\text{disease}|\text{positive})$ using frequencies out of one million

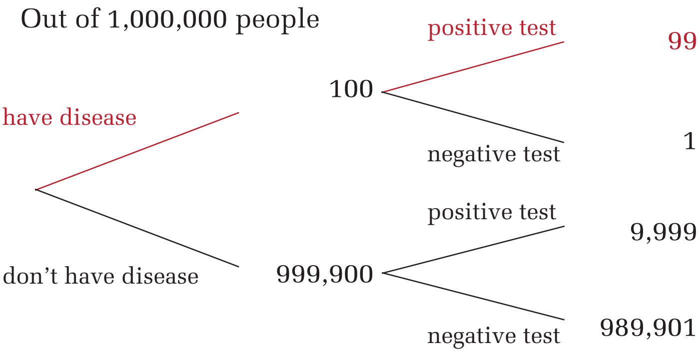

How many people out of 1,000,000 received positive diagnoses?

The red path above highlights one way: 99 people in 1,000,000 will get a positive diagnosis and have the disease

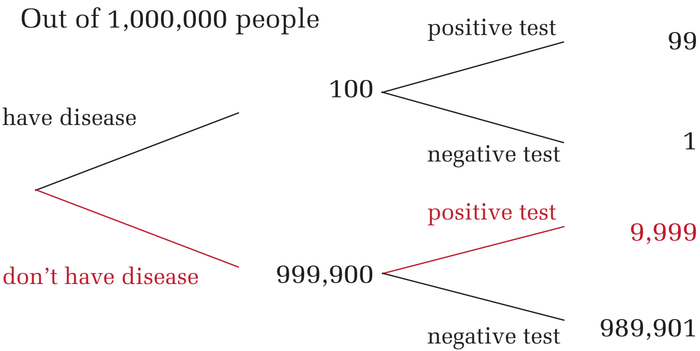

How many people out of 1,000,000 received positive diagnoses?

But there is another way: 9,999 people out of 1,000,000 will test positive even without the disease

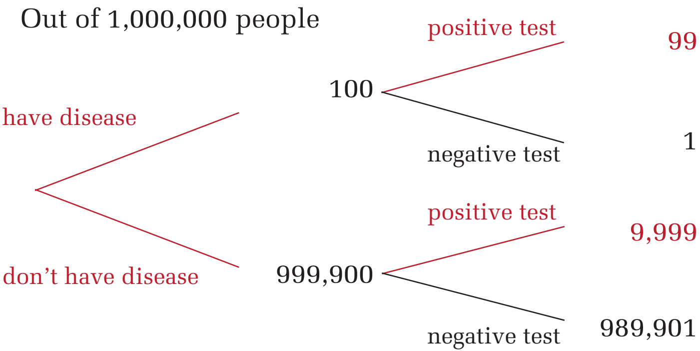

A total of $99 + 9,999$ receive positive diagnoses, but only 99 of these people have the disease.

$$\Pr(\text{disease}|\text{positive}) = \frac{99}{99 + 9,999} = 0.0098 \approx \text{ 1 percent probability}$$

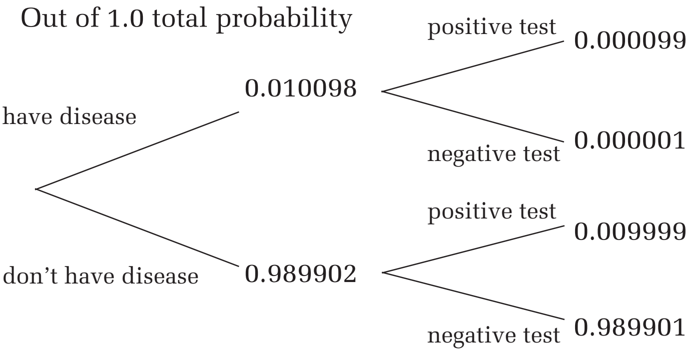

We could also solve this tree using the probabilities themselves, and the formula for conditional probability

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

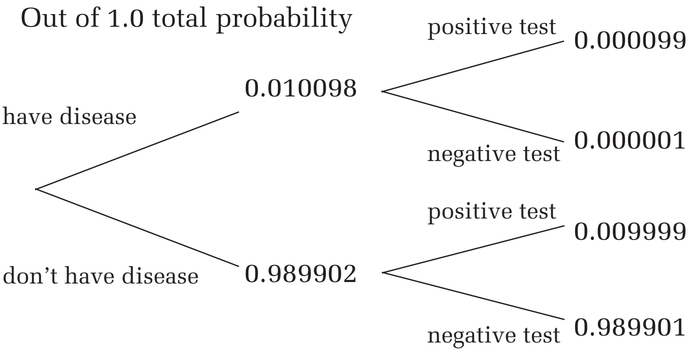

We could also solve this tree using the probabilities themselves, and the formula for conditional probability

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\Pr(\text{disease}|\text{positive}) = \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})}$$

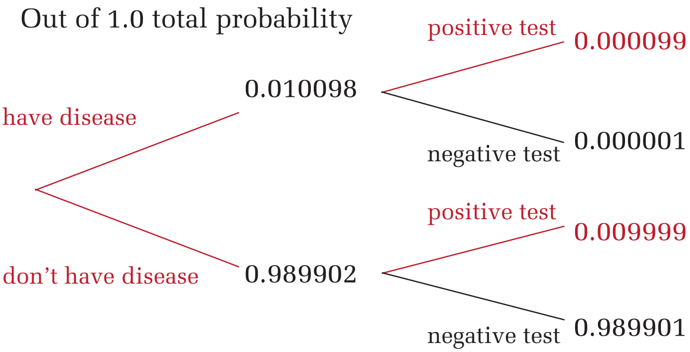

What is the marginal probability of a positive diagnosis, $\Pr(\text{positive})$?

The sum of the probabilities in red: $\Pr(\text{positive}) = 0.010098$

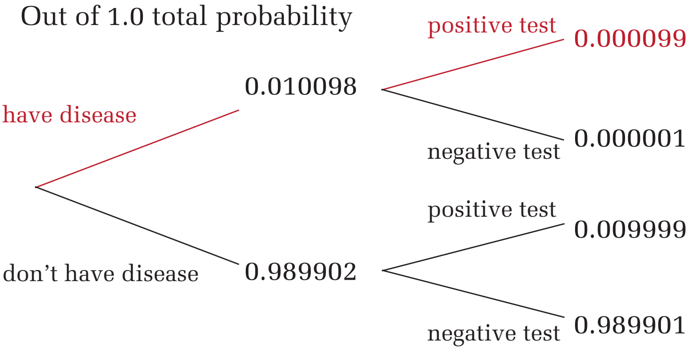

What is the joint probability of a positive diagnosis and a positive test?

The probability in red: $\Pr(\text{disease} \cap \text{positive}) = 0.000099$

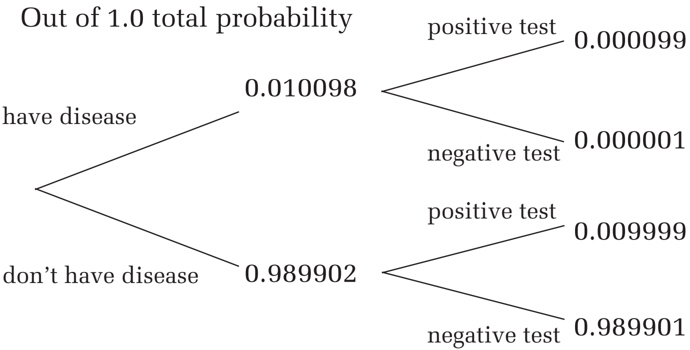

$$\Pr(\text{disease}|\text{positive}) = \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})}$$

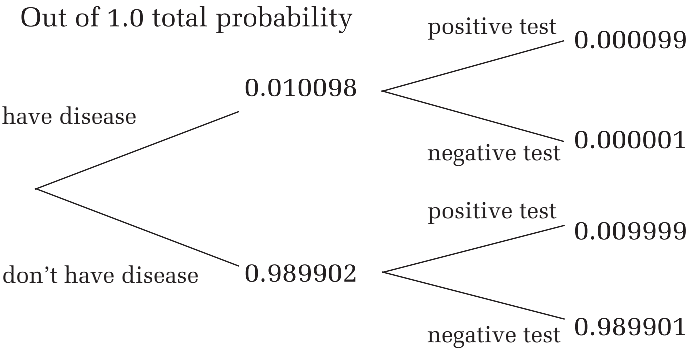

$$
\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease} \cap \text{positive})}{\Pr(\text{positive})} \\
&= \frac{0.000099}{0.010098} = 0.0098 \\
&\approx \text{1 percent probability}
\end{aligned}
$$

Things to ponder

1. It looks like a test for a rare disease would need to be staggeringly accurate before we trust it

   But remember we assumed the test was administered at random

   What if your doctor already suspected you had the disease.
   Would you still discount a positive result?

Things to ponder

1. It looks like a test for a rare disease would need to be staggeringly accurate before we trust it

   But remember we assumed the test was administered at random

   What if your doctor already suspected you had the disease. Would you still discount a positive result?

2. We can present probabilities as proportions (0.001) or as ratios (1 in 1000). Which do you find easier to understand?

Things to ponder

1. It looks like a test for a rare disease would need to be staggeringly accurate before we trust it

   But remember we assumed the test was administered at random

   What if your doctor already suspected you had the disease. Would you still discount a positive result?

2. We can present probabilities as proportions (0.001) or as ratios (1 in 1000). Which do you find easier to understand?

   This example was administered as a test question to a panel of doctors

   Using proportions, *most* doctors got the answer badly wrong (Kahneman & Tversky)

   Using ratios to understand the problem, most got it right! (Gigerenzer)

A statistician plays the lottery

Most US states run lotteries involving a daily drawing of 6 numbered balls from an urn, without replacement

In WA, lottery players pick six unique integers from 1 to 49

A player who picks all 6 numbers correctly wins the jackpot

If no one picks all 6 numbers, the jackpot rolls over to the next drawing

In WA and other states, partial matches win smaller prizes;
we will neglect these to keep our example simple

A statistician plays the lottery

Questions a statistician asks about a lottery:

1. What is the probability of winning the jackpot from buying a single ticket?
2. What is the expected return on a single ticket?
3. Can I increase my expected return using a strategy?
4. Based on the above, should I play the lottery?

Probability of winning the jackpot

What is the sample space?

Probability of winning the jackpot

What is the sample space?

All possible combinations of 6 different numbers chosen at random between 1 and 49.

What is the probability of matching such a number?

Start by identifying the event of interest:

Probability of winning the jackpot

What is the sample space?

All possible combinations of 6 different numbers chosen at random between 1 and 49.

What is the probability of matching such a number?

Start by identifying the event of interest:

Matching all 6 selected numbers

This is a complex event consisting of six sequential sub-events

Let's denote a successful match of the nth number as m_n

Probability of winning the jackpot

To have a “success”, we must see the following six events in order:

- Match the first draw (a random selection out of 49 numbers) to any of our 6 picks.

  We have six chances in 49 to get this right, so our probability is:

$$\Pr(m_1) = 6/49 = 0.1224$$

Probability of winning the jackpot

To have a "success", we must see the following six events in order:

- Match the first draw (a random selection out of 49 numbers) to any of our 6 picks.

  We have six chances in 49 to get this right, so our probability is:

  $$\Pr(m_1) = 6/49 = 0.1224$$

- Match the second draw (randomly chosen from 48 remaining numbers) to any of our remaining five picks.

  Assuming we got the first number, we have five chances in 48 to get this right:

  $$\Pr(m_2|m_1) = 5/48 = 0.1042$$

Probability of winning the jackpot

- Match the third draw (randomly chosen from 47 remaining numbers) to any of our remaining four picks.

  Assuming we got the first and second numbers, we have four chances in 47 to get this right:

$$\Pr(m_3|m_1, m_2) = 4/47 = 0.0851$$

Probability of winning the jackpot

- Match the third draw (randomly chosen from 47 remaining numbers) to any of our remaining four picks.

  Assuming we got the first and second numbers, we have four chances in 47 to get this right:

  $$\Pr(m_3|m_1, m_2) = 4/47 = 0.0851$$

- Match the fourth draw (randomly chosen from 46 remaining numbers) to any of our remaining three picks.

  Assuming we got the first, second, and third numbers, we have three chances in 46 to get this right:

  $$\Pr(m_4|m_1, m_2, m_3) = 3/46 = 0.0652$$

Probability of winning the jackpot

- Match the fifth draw (randomly chosen from 45 remaining numbers) to any of our remaining two picks.

  Assuming we got the first, second, third, and fourth numbers, we have two chances in 45 to get this right:

$$\Pr(m_5|m_1, m_2, m_3, m_4) = 2/45 = 0.0444$$

Probability of winning the jackpot

- Match the fifth draw (randomly chosen from 45 remaining numbers) to any of our remaining two picks.

  Assuming we got the first, second, third, and fourth numbers, we have two chances in 45 to get this right:

  $$\Pr(m_5|m_1, m_2, m_3, m_4) = 2/45 = 0.0444$$

- Match the sixth draw (randomly chosen from 44 remaining numbers) to our only remaining pick.

  Assuming we got the first, second, third, fourth, and fifth numbers, we have one chance in 44 to get this right:

  $$\Pr(m_6|m_1, m_2, m_3, m_4, m_5) = 1/44 = 0.0227$$

Probability of winning the jackpot

Note something interesting: matching numbers gets harder as we go, because our set of possible matches is shrinking.

To win the jackpot, we must match all six numbers,
so we need the joint probability of the above events.

We use the general rule for calculating the joint probability of multiple events:

$$\Pr(\text{Jackpot}) = \Pr(m_1 \cap m_2 \cap m_3 \cap m_4 \cap m_5 \cap m_6)$$

Probability of winning the jackpot

Note something interesting: matching numbers gets harder as we go, because our set of possible matches is shrinking.

To win the jackpot, we must match all six numbers,
so we need the joint probability of the above events.

We use the general rule for calculating the joint probability of multiple events:

$$\begin{aligned}
\Pr(\text{Jackpot}) &= \Pr(m_1 \cap m_2 \cap m_3 \cap m_4 \cap m_5 \cap m_6) \\
&= \Pr(m_1) \times \Pr(m_2|m_1) \times \Pr(m_3|m_1, m_2) \times \Pr(m_4|m_1, m_2, m_3) \\
&\qquad \times \Pr(m_5|m_1, m_2, m_3, m_4) \times \Pr(m_6|m_1, m_2, m_3, m_4, m_5)
\end{aligned}$$

Probability of winning the jackpot

Note something interesting: matching numbers gets harder as we go, because our set of possible matches is shrinking.

To win the jackpot, we must match all six numbers, so we need the joint probability of the above events.

We use the general rule for calculating the joint probability of multiple events:

$$\begin{aligned}
\Pr(\text{Jackpot}) &= \Pr(m_1 \cap m_2 \cap m_3 \cap m_4 \cap m_5 \cap m_6) \\
&= \Pr(m_1) \times \Pr(m_2|m_1) \times \Pr(m_3|m_1, m_2) \times \Pr(m_4|m_1, m_2, m_3) \\
&\qquad \times \Pr(m_5|m_1, m_2, m_3, m_4) \times \Pr(m_6|m_1, m_2, m_3, m_4, m_5) \\
&= \frac{6}{49} \times \frac{5}{48} \times \frac{4}{47} \times \frac{3}{46} \times \frac{2}{45} \times \frac{1}{44}
\end{aligned}$$

Probability of winning the jackpot

Note something interesting: matching numbers gets harder as we go, because our set of possible matches is shrinking.

To win the jackpot, we must match all six numbers, so we need the joint probability of the above events.

We use the general rule for calculating the joint probability of multiple events:

$$\begin{aligned}
\Pr(\text{Jackpot}) &= \Pr(m_1 \cap m_2 \cap m_3 \cap m_4 \cap m_5 \cap m_6) \\
&= \Pr(m_1) \times \Pr(m_2|m_1) \times \Pr(m_3|m_1, m_2) \times \Pr(m_4|m_1, m_2, m_3) \\
&\qquad \times \Pr(m_5|m_1, m_2, m_3, m_4) \times \Pr(m_6|m_1, m_2, m_3, m_4, m_5) \\
&= \frac{6}{49} \times \frac{5}{48} \times \frac{4}{47} \times \frac{3}{46} \times \frac{2}{45} \times \frac{1}{44} \\
&= \frac{1}{13,983,816} = 0.0000000715
\end{aligned}$$

What is the expected return from a single ticket?

Suppose a lottery ticket costs one dollar

How much, on average, you expect to get back from the dollar you spend on a lottery ticket is the *expected value* of the ticket.

Let's denote the specific set of six number we chose as ticket_i, and the amount of money in the jackpot as J.

$$\text{E}(\text{ticket}_i) = J \times \text{Pr}(\text{ticket}_i \text{ wins})$$

Is the lottery a good investment?

Suppose the jackpot is 1 million dollars. Then

$$\mathrm{E}(\mathrm{ticket}_i) = 1,000,000 \times 0.0000000715 = 0.07$$

If you played the lottery many millions of times at a $1 million jackpot, you would expect to get back an average of 7 cents for every dollar in tickets purchased

To break even, the jackpot would have to be about $14 million each time

But even then, over any short run of lotteries, you would expect nothing.

The lottery is high risk, with most of the expected return coming from a low probability event

What is the expected return from a single ticket? (redux)

But wait! There are other players, and nothing prevents them from picking the same six numbers!

If we pick the winning numbers, but so do q other people, we will have to split the jackpot $q+1$ ways!

And the bigger the jackpot, the more people play...

So really,

$$\begin{aligned} \mathrm{E}(\text{ticket}_i) \quad = \quad J \times & \frac{\Pr(\text{ticket}_i \text{ wins})}{[1 \times \Pr(\text{no one else chooses ticket}_i)} \\ & +2 \times \Pr(\text{one other player chooses ticket}_i) \\ & \quad +3 \times \Pr(\text{two other players choose ticket}_i) \ldots] \end{aligned}$$

Split jackpots shrink our winnings a lot!

A strategy for the lottery

All lottery numbers are equally likely to appear.

No six numbers are more likely to win than any other.

So is there any strategy we can use to maximize our winnings?

A strategy for the lottery

All lottery numbers are equally likely to appear.

No six numbers are more likely to win than any other.

So is there any strategy we can use to maximize our winnings?

All numbers are equally likely to win,
but not all numbers are equally likely to split the jackpot

If we pick numbers no one else plays,
we win just as often as ever, but never have to split our winnings

A strategy for the lottery

Most people misunderstand the concept of a random number sequence.

The sequence

$$9, 15, 17, 20, 35, 37$$

is just as likely to appear as

$$1, 2, 3, 4, 5, 6$$

or

$$44, 45, 46, 47, 48, 49$$

But most lottery players would see these sequences as “unlikely”

(How would you minimize returns from the lottery?

A strategy for the lottery

Most people misunderstand the concept of a random number sequence.

The sequence

$$9, 15, 17, 20, 35, 37$$

is just as likely to appear as

$$1, 2, 3, 4, 5, 6$$

or

$$44, 45, 46, 47, 48, 49$$

But most lottery players would see these sequences as “unlikely”

(How would you minimize returns from the lottery? Play numbers appearing in recent birthdays and anniversaries!)

Is the lottery worth playing?

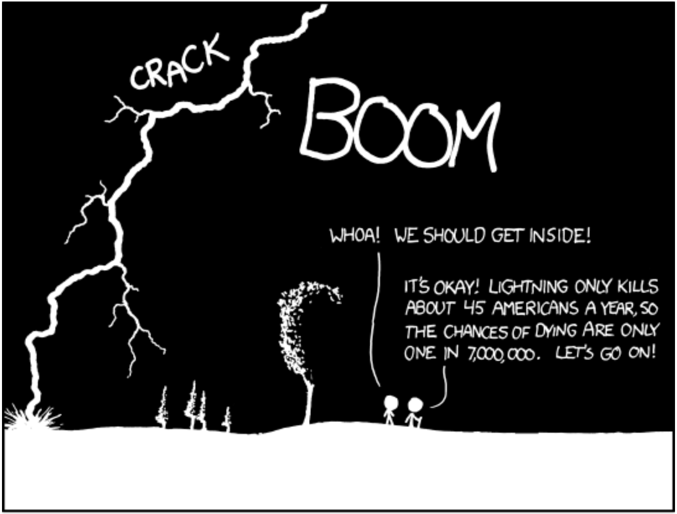

THE ANNUAL DEATH RATE AMONG PEOPLE WHO KNOW THAT STATISTIC IS ONE IN SIX.

http://xkcd.com/795/

The prosecutor's fallacy

DNA evidence can help clinch the case against a criminal suspect

Suppose the probability of a DNA match between two randomly selected people is 1 in 1,000,000

If a suspect's DNA matches evidence from a crime scene, guilt is very likely

But what if a prosecutor doesn't have a suspect?

The scenario

A prosecutor needs a suspect for a murder,
but has nothing but DNA from the crime scene

The prosecutor thinks the murderer could be anyone in the city,
which has 6 million residents

He begins testing them at random – and his first DNA match is you

Based on nothing but the DNA evidence, he charges you with the crime

He tells the grand jury that your guilt is beyond doubt,
because the probability of a DNA match by chance is just 1 in 1 million

You know you are innocent.
But how do you argue against the prosecutor's statistical argument?

But I'm innocent!

What probability do we want to calculate?

But I'm innocent!

What probability do we want to calculate?

$$\Pr(\text{innocent}|\text{DNA match})$$

But I'm innocent!

What probability do we want to calculate?

$$\Pr(\text{innocent}|\text{DNA match})$$

What probability has the prosecutor calculated?

$$\Pr(\text{DNA match}|\text{innocent})$$

The prosecutor has confused one conditional probability with its inverse!

What probabilities do we know?

What is the probability that you are innocent, before we do the DNA test?

$$\Pr(\text{innocent}) = 5,999,999 \text{ in } 6,000,000 = 0.9999998$$

What probabilities do we know?

What is the probability that you are innocent, before we do the DNA test?

$$\Pr(\text{innocent}) = 5,999,999 \text{ in } 6,000,000 = 0.9999998$$

Suppose that there are exactly 6 matches to the DNA in the city.
Then what is the probability that you are innocent after a successful DNA match?

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{innocent} \cap \text{DNA match})}{\Pr(\text{DNA match})}$$

What probabilities do we know?

What is the probability that you are innocent, before we do the DNA test?

$$\Pr(\text{innocent}) = 5,999,999 \text{ in } 6,000,000 = 0.9999998$$

Suppose that there are exactly 6 matches to the DNA in the city.
Then what is the probability that you are innocent after a successful DNA match?

$$\begin{aligned}\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{innocent} \cap \text{DNA match})}{\Pr(\text{DNA match})} \\ &= \frac{5 \text{ in } 6,000,000}{6 \text{ in } 6,000,000} = \frac{5}{6}\end{aligned}$$

Confusing one probability for its inverse could get you railroaded!

The prosecutor's fallacy

Mining a sufficiently large dataset for a coincidence *guarantees* that you find what you were looking for, so that the act of finding that coincidence proves nothing

The prosecutor's fallacy is well-known. A prosecutor who misuses statistical evidence in this way would risk serious sanctions

But wait a minute! DNA evidence gets used successfully in court all the time! What gives?

Using DNA to conduct a fishing expedition doesn't yield strong evidence.

But a DNA match on a single prior suspect is different.

To see this, we need one more bit of probability theory: Bayes Rule.

Bayes Rule

Suppose we know $\Pr(B|A)$, but want to know $\Pr(A|B)$ instead?
That is, we want to invert conditional probabilities:

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

Bayes Rule

Suppose we know $\Pr(B|A)$, but want to know $\Pr(A|B)$ instead?
That is, we want to invert conditional probabilities:

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

The formula in words:

$$\text{conditional probability of } A \text{ given } B \; = \frac{\text{conditional probability of } B \text{ given } A \; \times \; \text{marginal probability of } A}{\text{marginal probability of } B}$$

Bayes Rule

Suppose we know $\Pr(B|A)$, but want to know $\Pr(A|B)$ instead?
That is, we want to invert conditional probabilities:

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

The formula in words:

$$\text{conditional probability of } A \text{ given } B \;=\; \frac{\text{conditional probability of } B \text{ given } A \;\times\; \text{marginal probability of } A}{\text{marginal probability of } B}$$

Special names for the elements of Bayes' Rule:

$$\text{posterior probability of } A \text{ given } B \;=\; \frac{\text{likelihood of } B \text{ given } A \;\times\; \text{prior probability of } A}{\text{prior probability of } B}$$

Bayes Theorem (derivation)

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

Bayes Theorem (derivation)

$$\begin{aligned}\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}}\\ \Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)}\end{aligned}$$

Bayes Theorem (derivation)

$$
\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \\
\Pr(B|A) &= \frac{\Pr(A \cap B)}{\Pr(A)}
\end{aligned}
$$

Bayes Theorem (derivation)

$$
\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \\
\Pr(B|A) &= \frac{\Pr(A \cap B)}{\Pr(A)} \\
\Pr(B|A)\Pr(A) &= \Pr(A \cap B)
\end{aligned}
$$

Bayes Theorem (derivation)

$$\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \\
\Pr(B|A) &= \frac{\Pr(A \cap B)}{\Pr(A)} \\
\Pr(B|A)\Pr(A) &= \Pr(A \cap B) \\
\Pr(A|B)\Pr(B) &= \Pr(A \cap B)
\end{aligned}$$

Bayes Theorem (derivation)

$$
\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \\
\Pr(B|A) &= \frac{\Pr(A \cap B)}{\Pr(A)} \\
\Pr(B|A)\Pr(A) &= \Pr(A \cap B) \\
\Pr(A|B)\Pr(B) &= \Pr(A \cap B) \\
\Pr(A|B)\Pr(B) &= \Pr(B|A)\Pr(A)
\end{aligned}
$$

Bayes Theorem (derivation)

$$
\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\Pr(A|B) &= \frac{\Pr(A \cap B)}{\Pr(B)} \\
\Pr(B|A) &= \frac{\Pr(A \cap B)}{\Pr(A)} \\
\Pr(B|A)\Pr(A) &= \Pr(A \cap B) \\
\Pr(A|B)\Pr(B) &= \Pr(A \cap B) \\
\Pr(A|B)\Pr(B) &= \Pr(B|A)\Pr(A) \\
\Pr(A|B) &= \frac{\Pr(B|A)\Pr(A)}{\Pr(B)}
\end{aligned}
$$

Bayes Rule is the foundation for a major branch of statistics, called Bayesian statistics. UW a major center for Bayesian research.

Bayes Rule and a DNA match

We can use Bayes rule to figure out the probability that a suspect is guilty given a successful DNA match.

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match})}$$

Bayes Rule and a DNA match

We can use Bayes rule to figure out the probability that a suspect is guilty given a successful DNA match.

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match})}$$

We need to rewrite the denominator to sum up the probability of a DNA match in each scenario:

$$\begin{aligned}\Pr(\text{DNA match}) \quad = \quad & \Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) \\ & + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})\end{aligned}$$

Bayes Rule and a DNA match

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

Bayes Rule and a DNA match

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

This formula requires several pieces of information:

Prior probability of innocence, $\Pr(\textbf{innocent})$ Let's set this to 95%, in accordance with the principle of "innocent until proven guilty".

Bayes Rule and a DNA match

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

This formula requires several pieces of information:

Prior probability of innocence, $\Pr(\textbf{innocent})$ Let's set this to 95%, in accordance with the principle of "innocent until proven guilty".

Likelihood of DNA against an innocent party, $\Pr(\textbf{DNA match}|\textbf{innocent})$ This is the crux of the problem. Let's set this to 1 in 1 million initially.

Bayes Rule and a DNA match

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

This formula requires several pieces of information:

Prior probability of innocence, $\Pr(\textbf{innocent})$ Let's set this to 95%, in accordance with the principle of "innocent until proven guilty".

Likelihood of DNA against an innocent party, $\Pr(\textbf{DNA match}|\textbf{innocent})$ This is the crux of the problem. Let's set this to 1 in 1 million initially.

Probability of a DNA match given a guilty party, $\Pr(\textbf{DNA match}|\textbf{not innocent})$ Let's suppose the test performs perfectly when the party is guilty, so this is 1.

Bayes Rule and a DNA match

Supposing the DNA test really only goes against the innocent 1 in 1 million tries, we have:

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

Bayes Rule and a DNA match

Supposing the DNA test really only goes against the innocent 1 in 1 million tries, we have:

$$
\begin{aligned}
\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\begin{array}{l}\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) \\ \quad +\Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})\end{array}} \\
&= \frac{0.000001 \times 0.95}{0.000001 \times 0.95 + 1 \times 0.05}
\end{aligned}
$$

Bayes Rule and a DNA match

Supposing the DNA test really only goes against the innocent 1 in 1 million tries, we have:

$$\begin{aligned}
\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})} \\
&= \frac{0.000001 \times 0.95}{0.000001 \times 0.95 + 1 \times 0.05} \\
&= 0.00002
\end{aligned}$$

When DNA evidence weighs against a single suspect –
even one accorded a strong presumption of innocence –
the updated probability of innocence is microscopic

Bayes Rule and a DNA match

Suppose the prosecutor is on a fishing expedition in a large population

He waits until he has a match before going to court,
so the probability of a match – regardless of innocence or guilt – is 1.0!

$$\Pr(\text{innocent}|\text{DNA match}) = \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) + \Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})}$$

Bayes Rule and a DNA match

Suppose the prosecutor is on a fishing expedition in a large population

He waits until he has a match before going to court,
so the probability of a match – regardless of innocence or guilt – is 1.0!

$$\begin{aligned}
\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\begin{array}{l}\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) \\ \quad +\Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})\end{array}} \\
&= \frac{1 \times 0.95}{1 \times 0.95 + 1 \times 0.05}
\end{aligned}$$

Bayes Rule and a DNA match

Suppose the prosecutor is on a fishing expedition in a large population

He waits until he has a match before going to court,
so the probability of a match – regardless of innocence or guilt – is 1.0!

$$\begin{aligned}
\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\begin{array}{l}\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) \\ \quad +\Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})\end{array}} \\
&= \frac{1 \times 0.95}{1 \times 0.95 + 1 \times 0.05} \\
&= 0.95
\end{aligned}$$

Bayes Rule tells us that in this case,
DNA evidence adds no evidence of guilt whatsoever

Bayes Rule and a DNA match

Suppose the prosecutor is on a fishing expedition in a large population

He waits until he has a match before going to court,
so the probability of a match – regardless of innocence or guilt – is 1.0!

$$
\begin{aligned}
\Pr(\text{innocent}|\text{DNA match}) &= \frac{\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent})}{\begin{array}{l}\Pr(\text{DNA match}|\text{innocent}) \times \Pr(\text{innocent}) \\ \quad +\Pr(\text{DNA match}|\text{not innocent}) \times \Pr(\text{not innocent})\end{array}} \\
&= \frac{1 \times 0.95}{1 \times 0.95 + 1 \times 0.05} \\
&= 0.95
\end{aligned}
$$

Bayes Rule tells us that in this case,
DNA evidence adds no evidence of guilt whatsoever

We could tweak this by changing the prior or likelihood to include our assumption that someone in town must be guilty –
but that's not likely to fly in court or get us very far

Let’s Make a Deal

On *Let’s Make a Deal*, host Monty Hall offers you the following choice:

1. There are 3 doors. Behind one is a car. Behind the other two are goats.
2. You choose a door. It stays closed.
3. Monty picks one of the two remaining doors, and opens it to reveal a goat.
4. Your choice: Keep the door you chose in step 1, or switch to the third door.

What should you do?

Let's Make a Deal

What is the probability problem here?

Let's Make a Deal

What is the probability problem here?

1. What is the probability of a car from staying?
2. What is the probability of a car from switching?
3. Which is bigger?

How can we solve the problem?

Let's Make a Deal

What is the probability problem here?

1. What is the probability of a car from staying?
2. What is the probability of a car from switching?
3. Which is bigger?

How can we solve the problem?

1. Use probability theory: Bayes' Rule
2. Use brute force: Monte Carlo simulation

Using probability theory can get hard for complex scenarios

Monte Carlo is equally easy no matter how complex the scenario, but requires programming

Bayes Rule Solution (1)

We have doors A, B, and C.

Ex ante, the probability the car is behind each of these doors is just

$$\Pr(A) = \Pr(B) = \Pr(C) = \frac{1}{3}$$

Since the contestant picks a door $\mathcal{D}$ at random,

$$\Pr(\mathcal{D}) = \frac{1}{3}$$

For the sake of argument, suppose the contestant chooses $\mathcal{D} = A$.

Bayes Rule Solution (2)

Now, Monty Hall picks a door $\mathcal{E}$ to show a goat.

We want to know the probability that the remaining door $\mathcal{F}$ hides a car *given Monty's exposure of door* $\mathcal{E}$.

For the sake of argument, suppose that Monty picks $\mathcal{E} = B$.

We need to calculate $\Pr(\mathcal{F}|\mathcal{E} = B)$. Use Bayes' Rule:

$$\Pr(a|b) = \frac{\Pr(b|a) \times \Pr(a)}{\Pr(b)}$$

or in our case,

$$\Pr(\mathcal{F}|\mathcal{E} = B) = \frac{\Pr(\mathcal{E} = B|\mathcal{F}) \times \Pr(\mathcal{F})}{\Pr(\mathcal{E} = B)}$$

Bayes Rule Solution (3)

The problem we need to solve:

$$\Pr(\mathcal{F}|\mathcal{E}=B) = \frac{\Pr(\mathcal{E}=B|\mathcal{F}) \times \Pr(\mathcal{F})}{\Pr(\mathcal{E}=B)}$$

We need $\Pr(\mathcal{E}=B|\mathcal{F})$, the probability Monty would open door B if the remaining door $\mathcal{F}=C$ actual held the car. By the rules of the game, this must be 1: Monty never shows the car.

We have the marginal probability that $\mathcal{F}=C$ holds the car; it's $1/3$.

Finally, we have the probability that Monty would "choose" to open B rather than C. This, of course, is $1/2$.

Substituting into Bayes' Rule, we find

$$\Pr(\mathcal{F}|\mathcal{E}=B) = \frac{1 \times \frac{1}{3}}{\frac{1}{2}} = \frac{2}{3}$$

Monty Hall: Intuitive solution

Can we explain the Monty Hall problem without explaining Bayes Rule?

Key is to notice that Monty is *filtering out a goat* from the two remaining doors:

- Assume probability of a car behind each door is 1/3, ex ante
- Collectively, probability of car behind door 2 *or* 3 is 2/3
- Monty revealing a goat shows us where car must be if it was behind 2 or 3
- In effect, Monty is giving us a choice to take *any* car behind door 1 *or any* car behind doors 2 *and* 3

Solving Monty Hall without math or intuition

What if we can't wrap our brains around the intuition of a complex random situation, or figure out how to apply probability theory to it?

We can still calculate any probability we want using *simulation*

We could do this with scraps of paper and a hat:
just act out the random experiment, and see what happens

Only problem is we'll need to repeat the experiment 1000s of times to estimate probabilities precisely

So let's do it on the computer instead!

Warning: programming content

The last part of this lecture is a bit advanced

Shows what you can do in `R` with a little programming knowledge

Not required that you learn how to do this for 321.
But it may be useful later in your academic or professional careers

Pseudo-code: A sketch of the desired R program

1. Set up the doors, goats, and car
2. Contestant picks a door
3. Monty “picks” a remaining door
4. Record where the car and goats were
5. Do all of the above many many times
6. Print the fraction of times a car was found

```
# Monty Hall Problem
# Chris Adolph
# 1/6/2005

sims <- 10000            # Simulations run
doors <- c(1,0,0)        # The car (1) and the goats (0)
cars.chosen <- 0         # Save cars from first choice here
cars.rejected <- 0       # Save cars from switching here

for (i in 1:sims) {      # Loop through simulations

  # First, contestant picks a door
  first.choice <- sample(doors,3,replace=F)

  # Choosing a door means rejecting the other two
  chosen <- first.choice[1]
  rejected <- first.choice[2:3]
```

```
  # Monty Hall removes a goat from the rejected doors
  rejected <- sort(rejected)
  if (rejected[1]==0)
    rejected <- rejected[2]

  # Record keeping:  where was the car?
  cars.chosen <- cars.chosen + chosen
  cars.rejected <- cars.rejected + rejected

}

# Print results
cat("Probability of a car from staying with 1st door",
    cars.chosen/sims,"\n")

cat("Probability of a car from switching to 2nd door",
    cars.rejected/sims,"\n")
```