4.5 — Bayesian Players

ECON 316 • Game Theory • Fall 2021

Ryan Safner

Assistant Professor of Economics

safner@hood.edu

ryansafner/gameF21

gameF21.classes.ryansafner.com

Outline

Bayesian Statistics

Bayes’ Rule Example

Bayesian Statistics

Bayesian Statistics

- Most people's understanding & intuitions of probability are about the **objective** **frequency** of events occurring
  - "If I flip a fair coin many times, the probability of Heads is 0.50"
  - "If this election were repeated many times, the probability of Biden winning is 0.60"
- This is known as the **"frequentist"** interpretation of probability
  - And is almost entirely the only thing taught to students (because it's easier to explain)

Bayesian Statistics

- Another valid (competing) interpretation is probability represents our **subjective belief** about an event
    - "I am 50% certain the next coin flip will be Heads"
    - "I am 60% certain that Biden will win the election"
    - This is particularly useful for **unique** events (that occur once…and really, isn't that every event in the real world?)
- This is known as the "Bayesian" interpretation of probability

Bayesian Statistics

- In Bayesian statistics, probability measures the degree of certainty about an event
  - Beliefs range from impossible $(p = 0)$ to certain $(p = 1)$
- This conditions probability on your **beliefs** about an event

Rev. Thomas Bayes

1702—1761

Bayesian Statistics

- The bread and butter of thinking like a Bayesian is **updating your beliefs in response to new evidence**
    - You have some **prior** belief about something
    - New evidence should **update** your belief (level of certainty) about it
    - Updated belief known as your **posterior** belief
- Your beliefs are *not* completelydetermined by the latest evidence, new evidence just *slightly* changes your beliefs, proportionate to how compelling the evidence is
- **This is fundamental to modern science and having rational beliefs**
    - And some mathematicians will tell you, the *proper* use of statistics

Bayesian Statistics Examples

1. You are a bartender. If the next person that walks in is wearing a kilt, what is the probability s/he wants to order Scotch?
2. You are playing poker and the player before you raises.
3. What is the probability that someone has watched the Superbowl? What if you learn that person is a man?
4. You are a policymaker deciding foreign policy, and get a new intelligence report.
5. You are trying to buy a home and make an offer, which the seller declines.

Conditional Probability

- All of this revolves around **conditional probability**: the probability of some event B occurring, given that event A has already occurred

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

- $P(B|A)$: “Probability of B given A”

Conditional Probability

- All of this revolves around **conditional probability**: the probability of some event B occurring, given that event A has already occurred

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

- If we know A has occurred, $P(A) > 0$, and then every outcome that is $\neg A$ ("not A") cannot occur $(P(\neg A) = 0)$

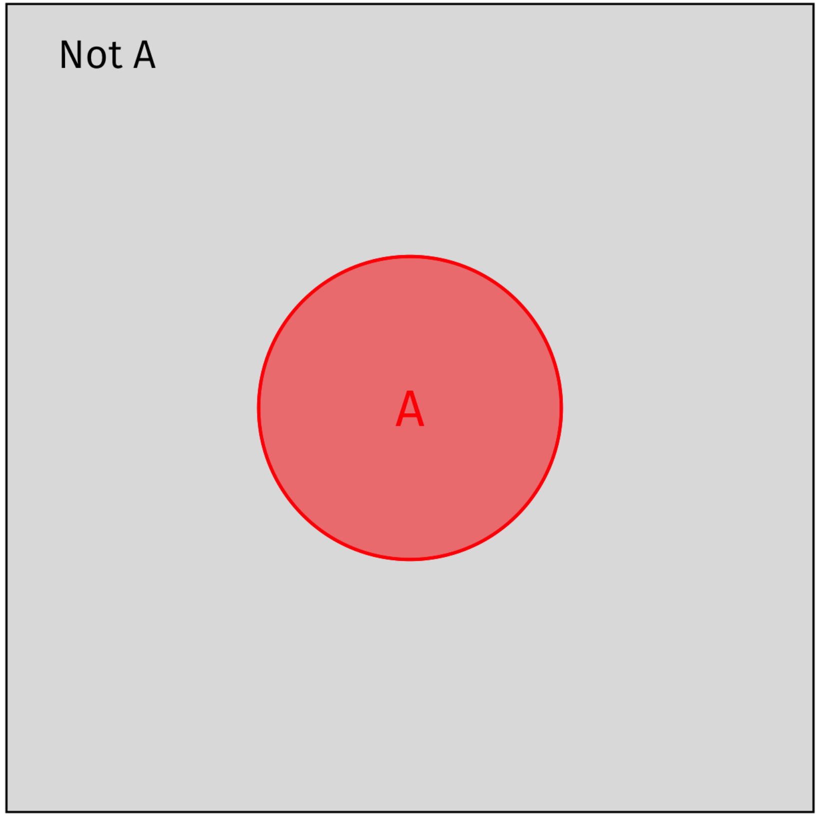

Conditional Probability

- All of this revolves around **conditional probability**: the probability of some event B occurring, given that event A has already occurred

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

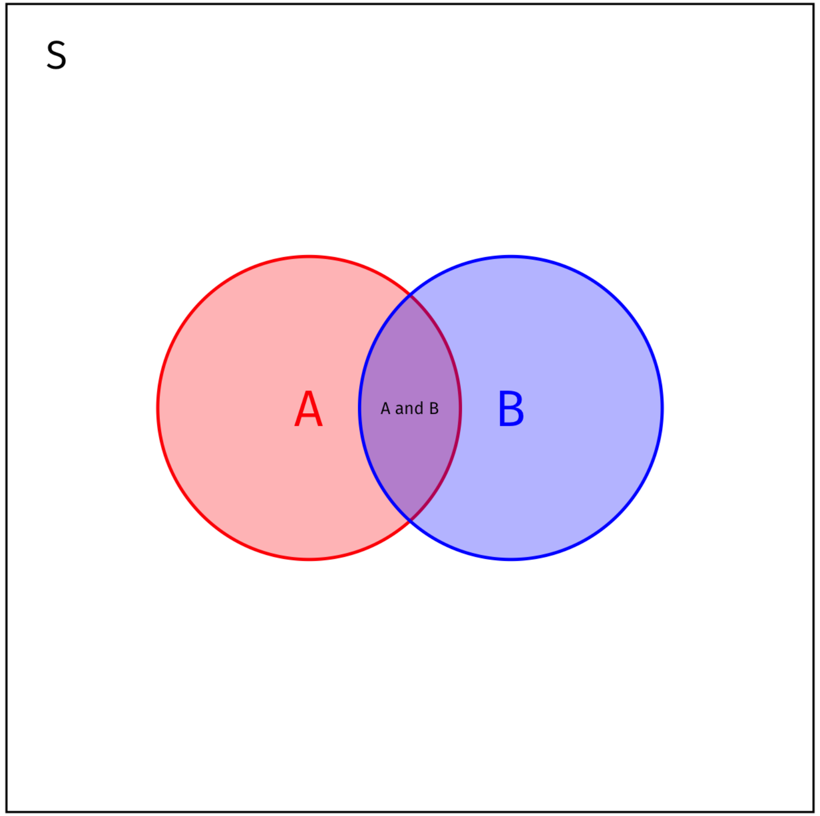

Conditional Probability

- All of this revolves around **conditional probability**: the probability of some event B occurring, given that event A has already occurred

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

- If we know A has occurred, $P(A) > 0$, and then every outcome that is $\neg A$ ("not A") cannot occur $(P(\neg A) = 0)$
  - The only part of B which can occur if A has occurred is A and B
  - Since the sample space S must equal 1, we've reduced the sample space to A, so we must rescale by $\frac{1}{P(A)}$

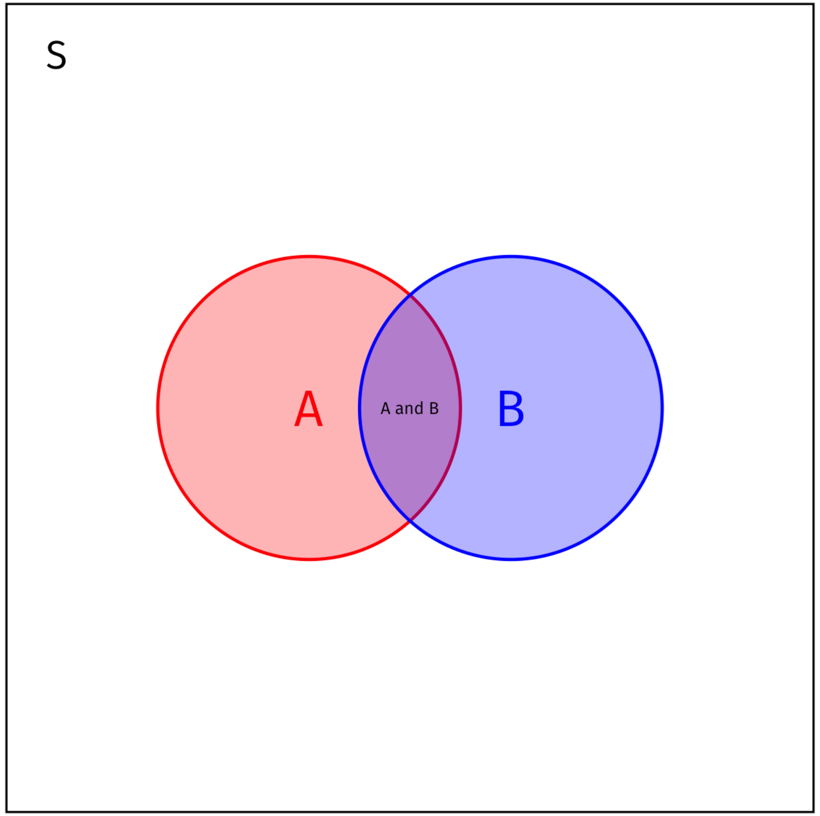

Conditional Probability

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

- If events A and B were **independent**, then the probability $P(A \text{ and } B)$ happening would be just $P(A) \times P(B)$
  - $P(A|B) = P(A)$
  - $P(B|A) = P(B)$

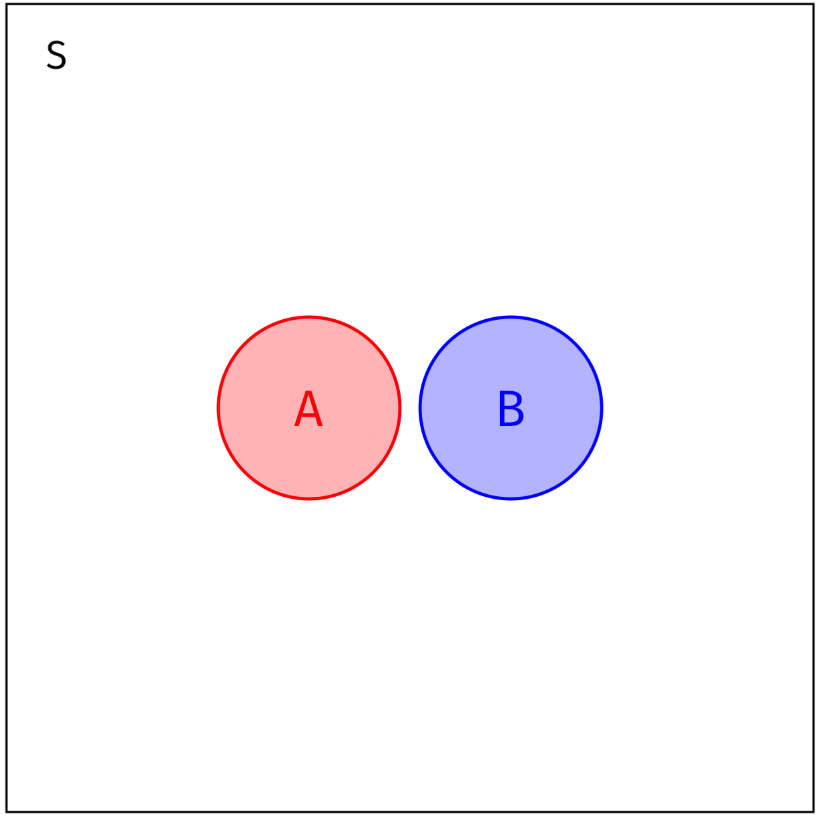

Conditional Probability

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

- But if they are *not* independent, it's $P(A \text{ and } B) = P(A) \times P(B|A)$
  - (Just multiplying both sides above by the denominator, $P(A)$)

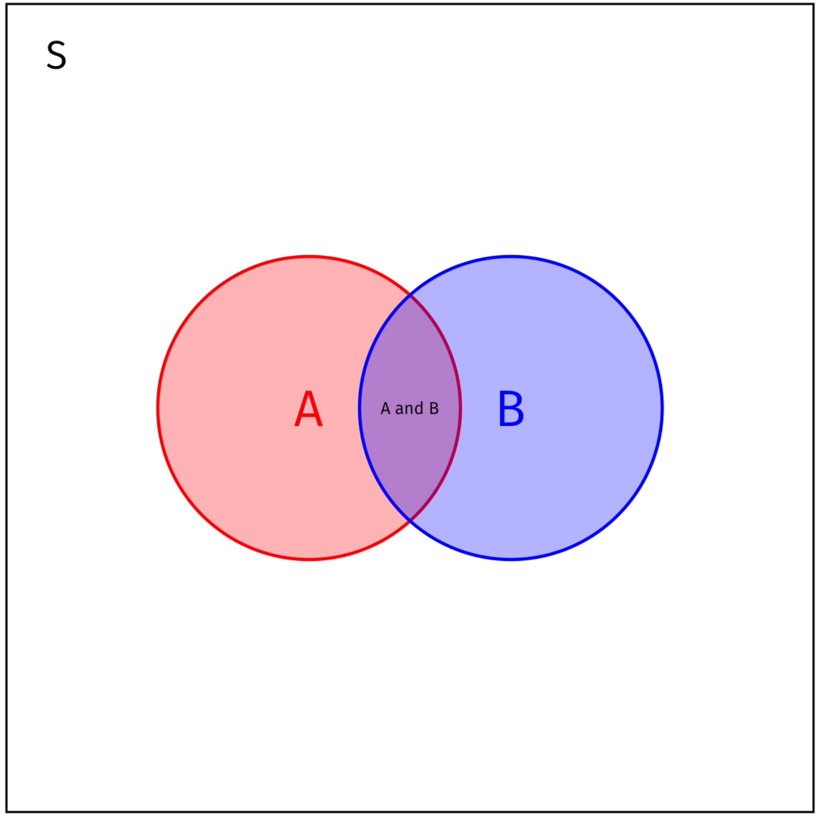

Conditional Probability and Bayes' Rule

- Bayes realized that the conditional probabilities of two non-independent events are proportionately related

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

$$P(A|B) = \frac{P(A \text{ and } B)}{P(B)}$$

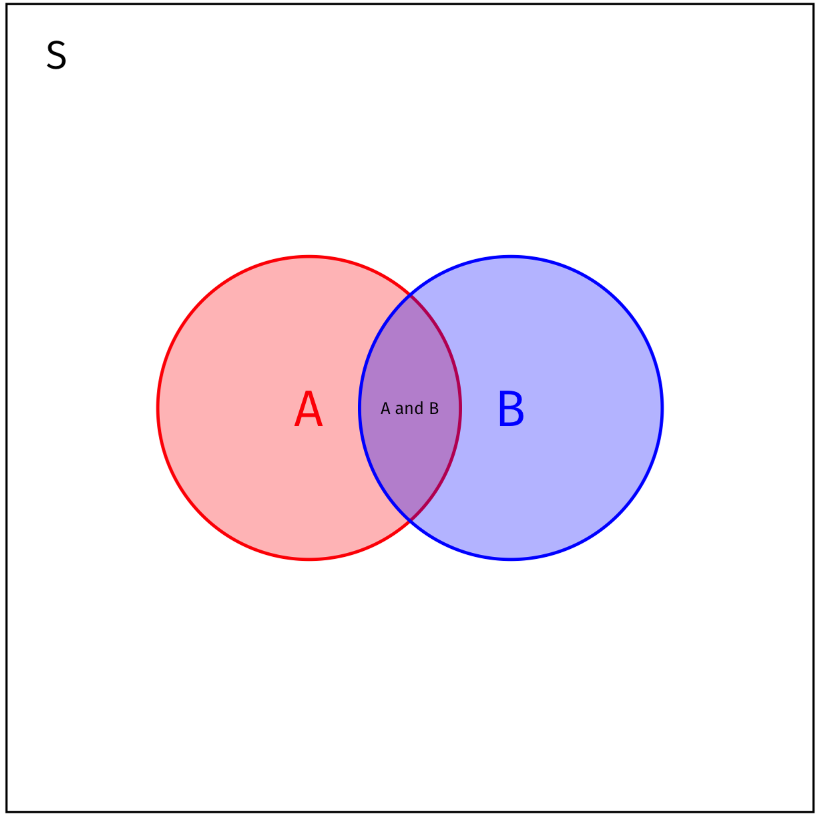

Conditional Probability and Bayes' Rule

- Bayes realized that the conditional probabilities of two non-independent events are proportionately related

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

$$P(A|B) = \frac{P(A \text{ and } B)}{P(B)}$$

$$P(A|B)P(B) = P(A \text{ and } B) = P(B|A)P(A)$$

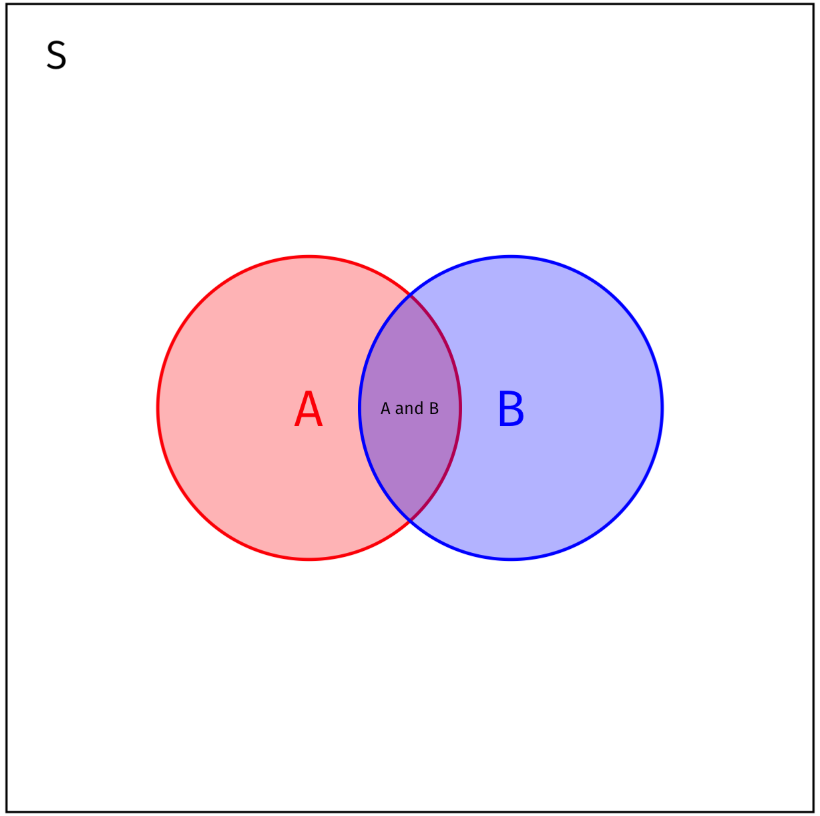

Conditional Probability and Bayes' Rule

- Bayes realized that the conditional probabilities of two non-independent events are proportionately related

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}$$

$$P(A|B) = \frac{P(A \text{ and } B)}{P(B)}$$

$$P(A|B)P(B) = P(A \text{ and } B) = P(B|A)P(A)$$

- Divide everything by P(B), you get, famously, **Bayes' rule**:

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

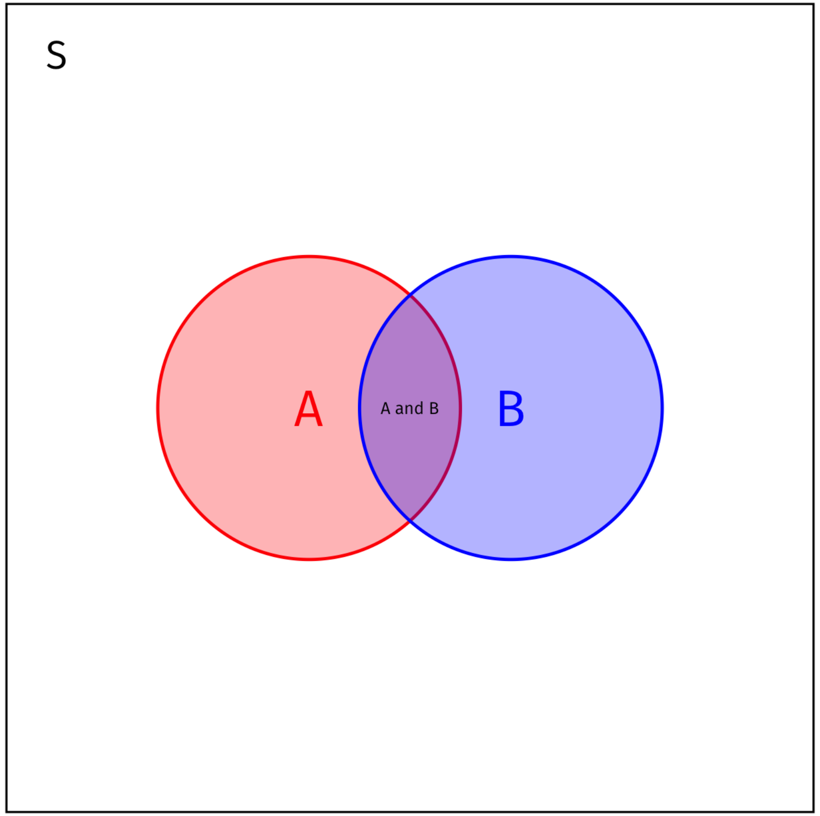

Bayes' Rule: Hypotheses and Evidence

- The A's and B's are rather difficult to remember if you don't use this often

- A lot of people prefer to think of **Bayes' rule** in terms of a hypothesis you have (H), and new evidence or data e

$$P(H|e) = \frac{P(e|H)p(H)}{P(e)}$$

- $P(H|e)$: **posterior** your hypothesis is correct given the new evidence
- $P(e|H)$: **likelihood** of seeing the evidence under your hypothesis
- $P(H)$: **prior belief** in of your hypothesis

Bayes' Rule Example

Bayes' Rule Example

Example: Suppose 1% of the population has a rare disease. A test that can diagnose the disease is 95% accurate. What is the probability that a person who takes the test and comes back positive has the disease?

- What would you guess the probability is?

Bayes' Rule Example

Example: Suppose 1% of the population has a rare disease. A test that can diagnose the disease is 95% accurate. What is the probability that a person who takes the test and comes back positive has the disease?

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95 = P(-|\neg\text{Disease})$
- We know $P(+|\text{Disease})$ but want to know $P(\text{Disease}|+)$
  - **These are not the same thing!**
  - Related by Bayes' Rule:

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

Bayes' Rule Example

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95 = P(-|\neg\text{Disease})$

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

- What is $P(+)$??

Bayes' Rule Example

- What is the total probability of B in the diagram?

$$P(B) = P(B \text{ and } A) + P(B \text{ and } \neg A)$$
$$= P(B|A)P(A) + P(B|\neg A)P(\neg A)$$

- This is known as the law of total probability

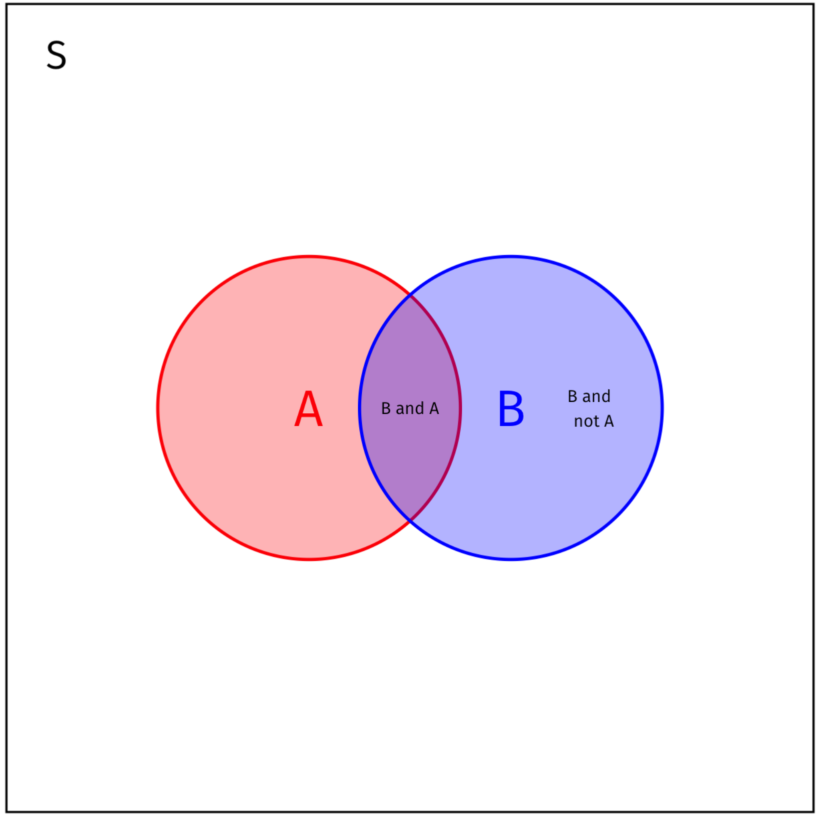

Bayes' Rule Example: Aside

- Because we usually have to figure out $P(B)$ (the denominator), Bayes' rule is often expanded to

$$P(B|A) = \frac{P(A|B)P(A)}{P(B|A)P(A) + P(B|\neg A)P(\neg A)}$$

- Assuming there are two possibilities (A and $\neg A$), e.g. True or False

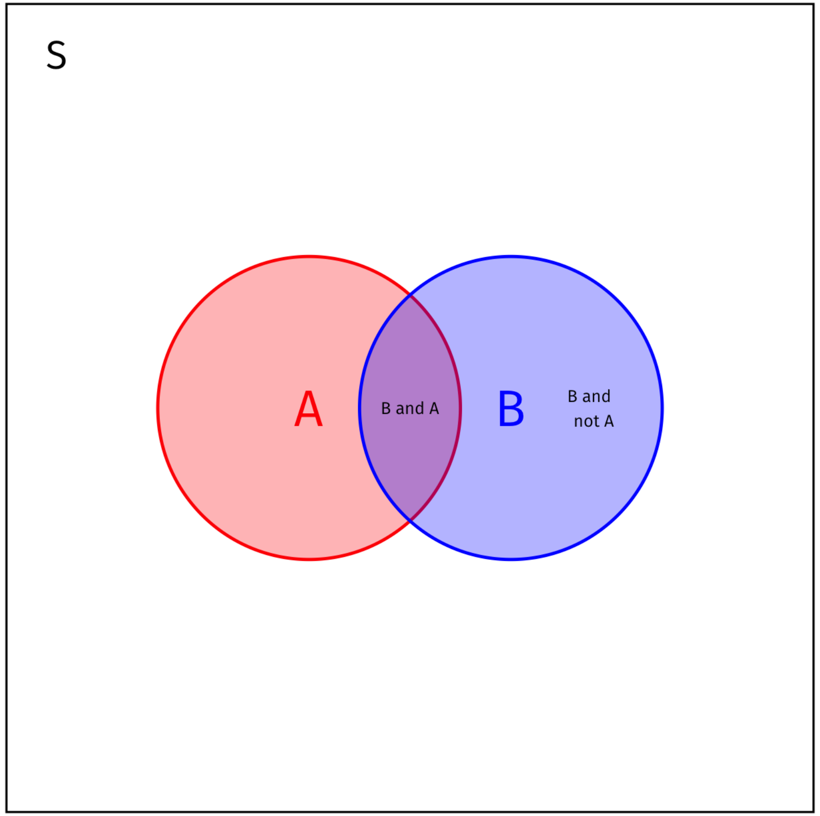

Bayes' Rule Example: Aside

- If there are more than two possibilities, you can further expand it to $\sum_{i=1}^{n} P(B|A_i)P(A_i)$ for n number of possible alternatives to A

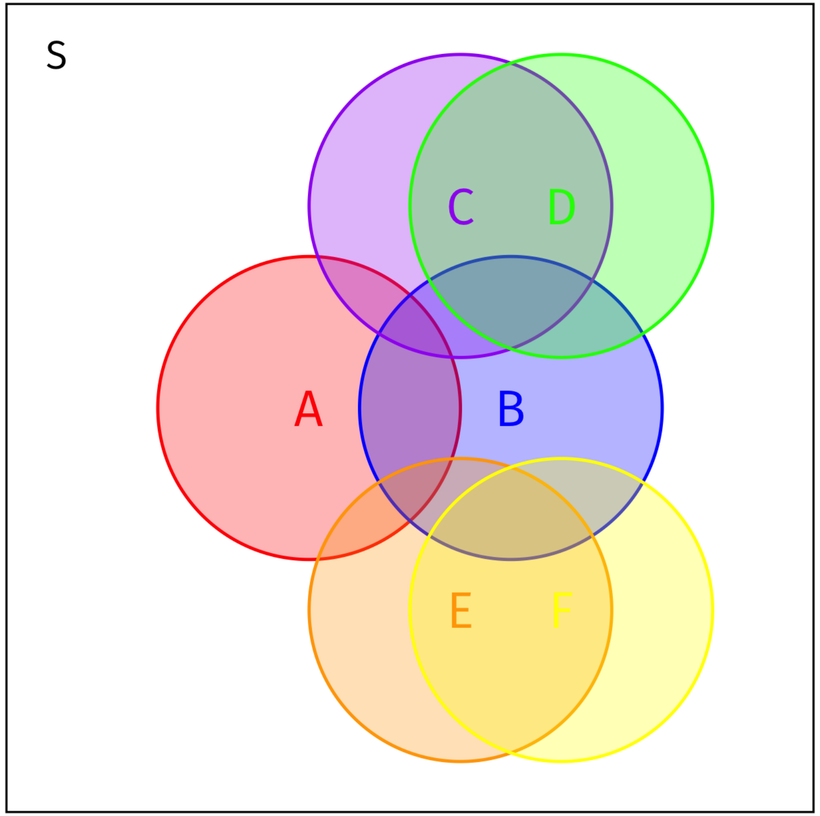

Bayes' Rule Example

- What is the total probability of $+$?

$$
\begin{aligned}
P(+) &= P(+ \text{ and Disease}) + P(+ \text{ and } \neg \text{ Disease}) \\
&= P(+|\text{Disease})P(\text{Disease}) + P(+|\neg\text{Disease})P(\neg\text{Disease})
\end{aligned}
$$

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95$

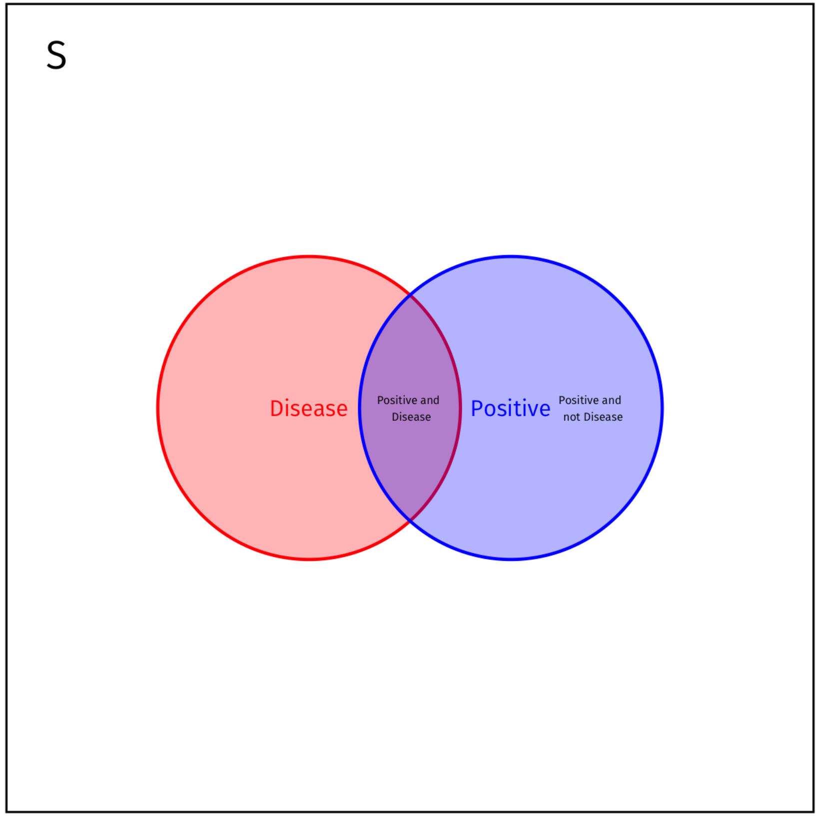

Bayes' Rule Example

- What is the total probability of +?

$$P(+) = P(+ \text{ and Disease}) + P(+ \text{ and} \neg \text{ Disease})$$
$$= P(+|\text{Disease})P(\text{Disease}) + P(+|\neg\text{Disease})P(\neg\text{Disease})$$

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95$

$$P(+) = 0.95(0.01) + 0.05(0.99) = 0.0590$$

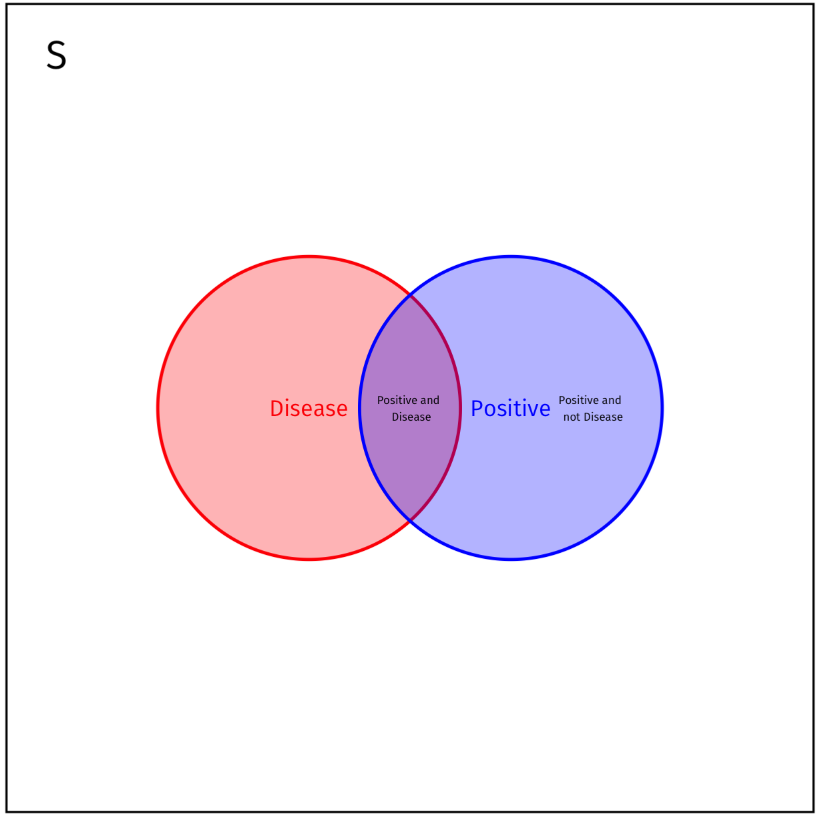

Bayes' Rule Example

	Disease	¬ Disease	Total
+	0.0095	0.0495	**0.0590**
-	0.0005	0.9405	0.9410
Total	0.0100	0.9900	1.0000

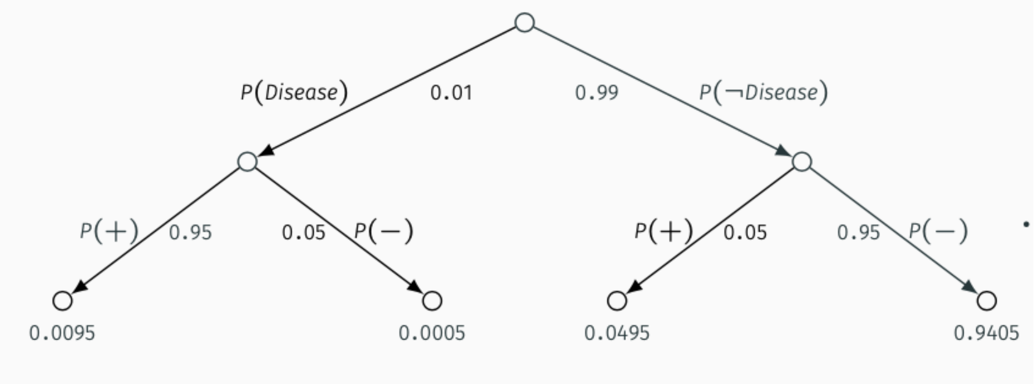

Bayes' Rule Example

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95 = P(-|\neg\text{Disease}) = 0.95$
- $P(+) = 0.0590$

Bayes’ Rule Example

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95 = P(-|\neg\text{Disease}) = 0.95$
- $P(+) = 0.0590$

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

Bayes' Rule Example

- $P(\text{Disease}) = 0.01$
- $P(+|\text{Disease}) = 0.95 = P(-|\neg\text{Disease}) = 0.95$
- $P(+) = 0.0590$

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

$$P(\text{Disease}|+) = \frac{0.95 \times 0.01}{0.0590}$$

$$= 0.16$$

- The probability you have the disease is only 16%!
    - Most people vastly overestimate because they forget the base rate of the disease, $P(\text{Disease})$ is so low (1%)!

Bayes' Rule and Bayesian Updating

- Bayes Rule tells us how we should update our beliefs given new evidence

Likelihood
How probable is the evidence given that our hypothesis is true?

Prior
How probable was our hypothesis *before* observing the evidence?

$$P(H \mid e) = \frac{P(e \mid H)\, P(H)}{P(e)}$$

Posterior
How probable is our hypothesis given the observed evidence?
(Not directly computable)

Marginal
How probable is the new evidence under all possible hypotheses?
$P(e) = \sum P(e \mid H_i)\, P(H_i)$

Highly, Highly Recommended

Highly, Highly Recommended