Bayesian Methods

Quantitative Understanding in Biology

1 October 2020

Lecture Notes by Jason Banfelder

Slide Compilation and Demonstratives by Ariana Clerkin

Introduction: Frequentist vs. Bayesian

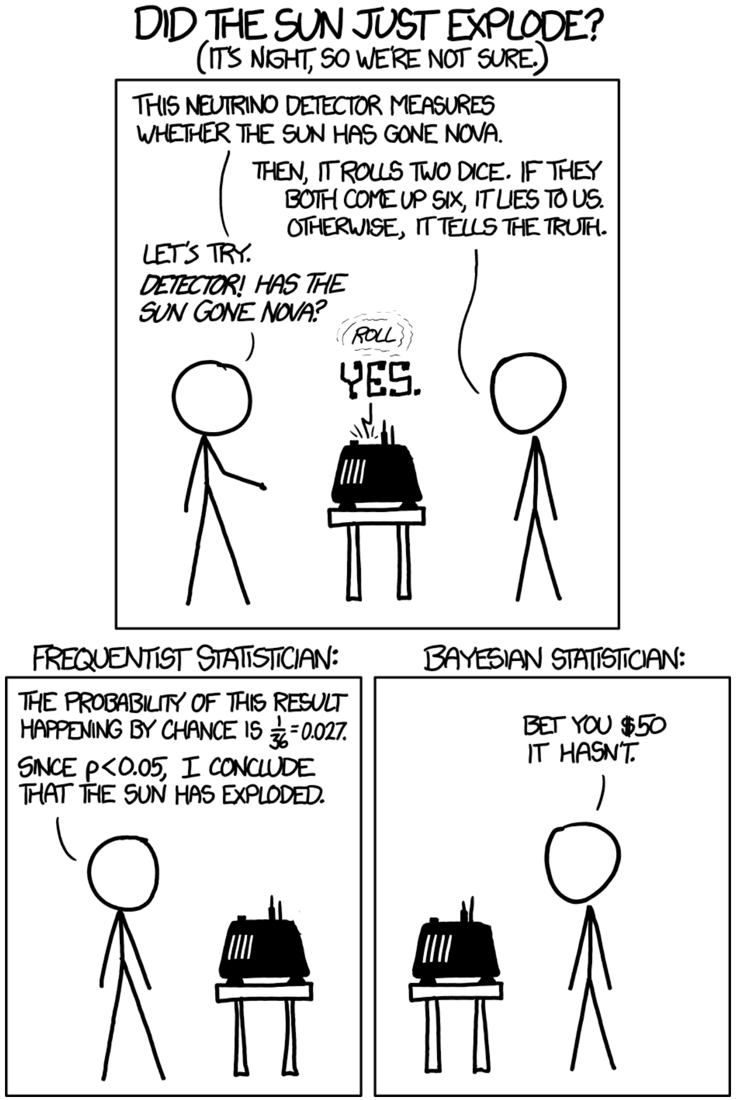

Source: https://xkcd.com/1133/

Probability
7th grade classroom

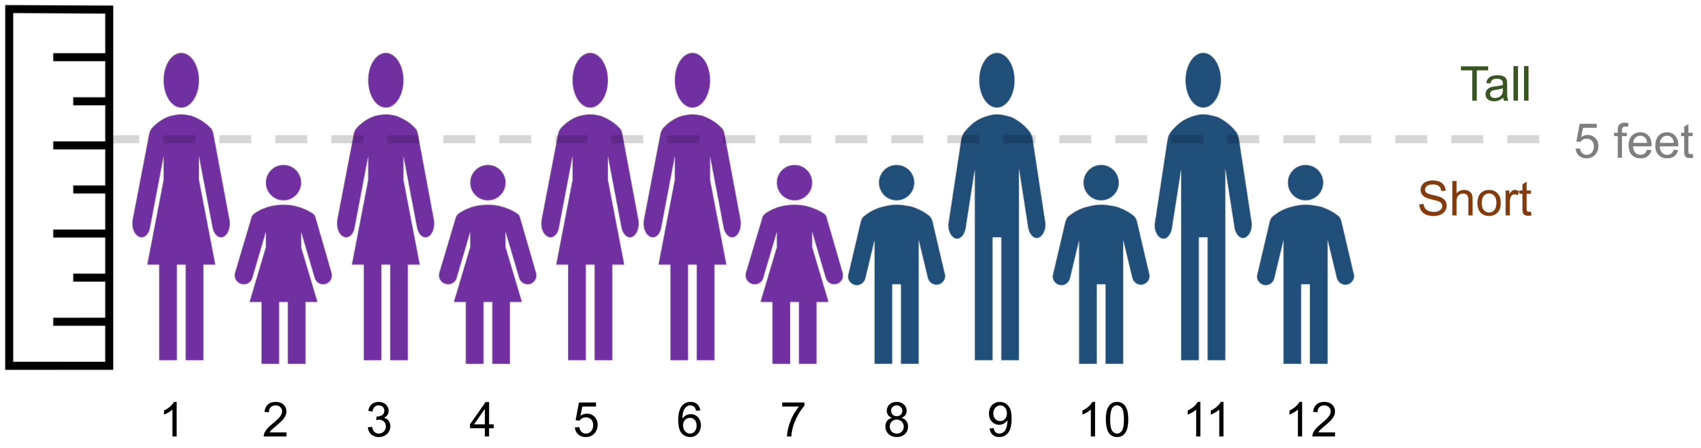

P(Tall) =

Conditional Probability
7^{th} Grade Classroom

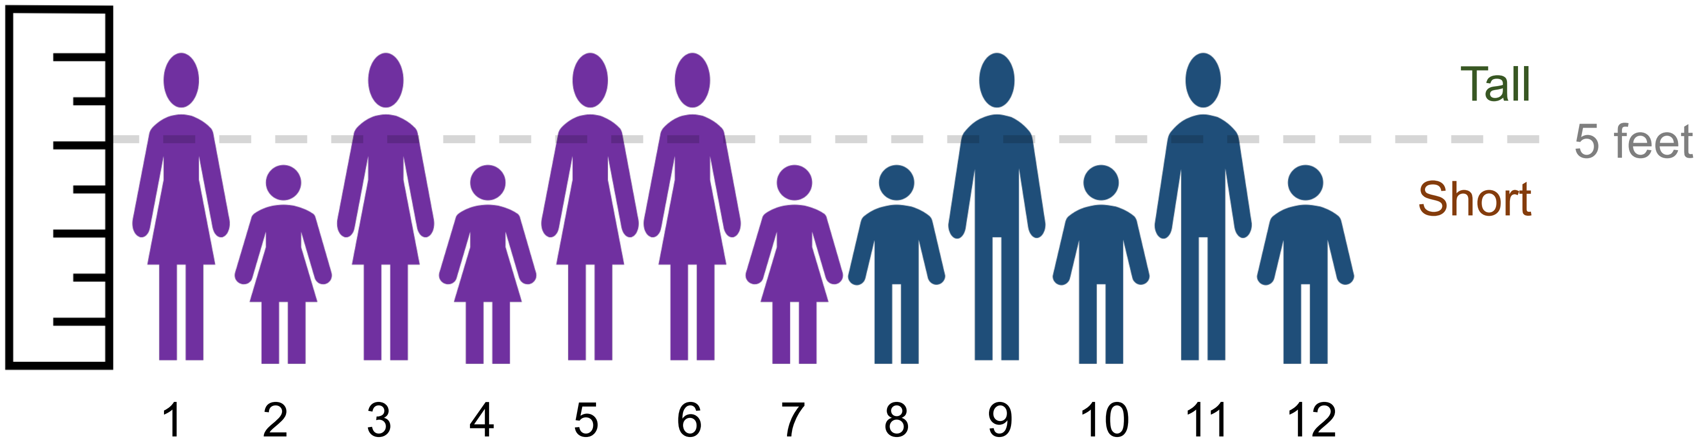

P(Tall | Female) = 4/7

Probability that the student is *Tall* *given that* the student is *Female* (Conditional Probability)

We expect P(Tall | Female) > P(Tall) without taking any measurements of this particular class.

This interplay goes both ways: generally P(Female | Tall) > P(Female)

Joint Probability

7^{th} Grade Classroom

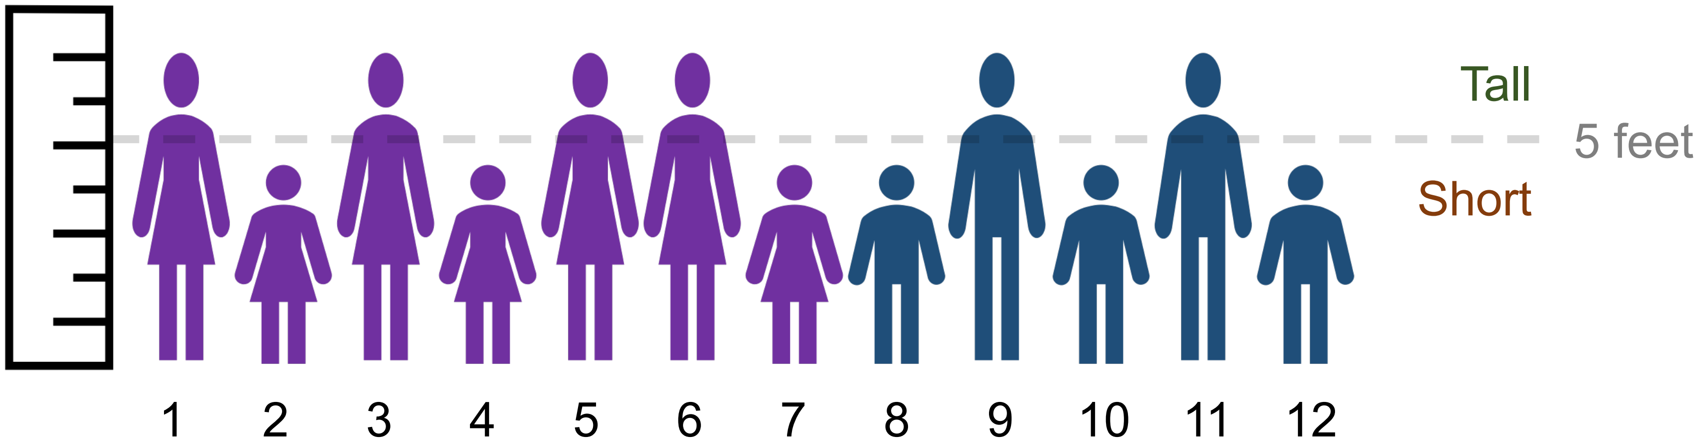

P(Tall , Female) = P(Female) · P(Tall | Female)

Probability that the student is *Tall and that* the student is *Female* (Joint Probability)

4/12 = 7/12 · 4/7

Joint Probability

7th Grade Classroom

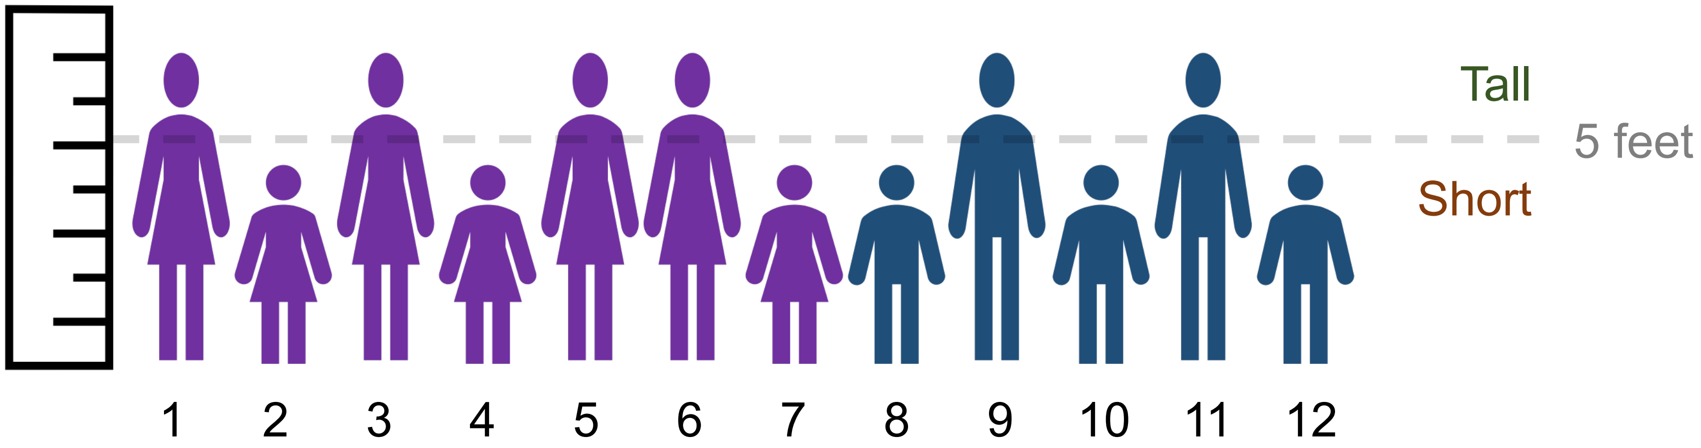

P(Tall , Female) = P(Female) · P(Tall | Female)

OR, equivalently

P(Tall , Female) = P(Tall) · P(Female | Tall)

4/12 = 6/12 · 4/6

Deriving Bayes' Rule

We have shown that:

$$P(\text{Tall}, \text{Female}) = P(\text{Female}) \cdot P(\text{Tall} \mid \text{Female})$$

$$P(\text{Tall}, \text{Female}) = P(\text{Tall}) \cdot P(\text{Female} \mid \text{Tall})$$

Therefore:

$$P(\text{Female}) \cdot P(\text{Tall} \mid \text{Female}) = P(\text{Tall}) \cdot P(\text{Female} \mid \text{Tall})$$

$$P(\text{Tall} \mid \text{Female}) = \frac{P(\text{Female} \mid \text{Tall}) \cdot P(\text{Tall})}{P(\text{Female})}$$

Or generally, for generic events A & B, we have

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

Bayes' Rule: Terminology

Likelihood

Prior Probability

Posterior Probability

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

Evidence or Marginal Likelihood

Applying Bayes' Rule

Information:

- 1% of women in a given population have breast cancer
- If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
- If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

Question:

What is the probability that a woman with a positive test result actually has cancer?

Multiple Choice:

Which notation shows the probability that a woman with a positive test result actually has cancer?

a.) P(Cancer | Positive Test)

b.) P(Cancer , Positive Test)

c.) P(Positive Test | Cancer)

d.) P(Positive Test ∩ Cancer)

Applying Bayes' Rule

- <u>Information:</u>
  - 1% of women in a given population have breast cancer
  - If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
  - If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

$$P(\text{Cancer} \mid \text{Positive}) = \frac{P(\text{Positive} \mid \text{Cancer}) \cdot P(\text{Cancer})}{P(\text{Positive})}$$

Applying Bayes' Rule

- Information:
  - 1% of women in a given population have breast cancer
  - If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
  - If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

$$P(\text{Cancer} \mid \text{Positive}) = \frac{P(\text{Positive} \mid \text{Cancer}) \cdot \overbrace{P(\text{Cancer})}^{0.01}}{P(\text{Positive})}$$

Applying Bayes' Rule

- Information:
  - 1% of women in a given population have breast cancer
  - If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
  - If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

$$P(\text{Cancer} \mid \text{Positive}) = \frac{\overbrace{P(\text{Positive} \mid \text{Cancer})}^{0.9} \cdot \overbrace{P(\text{Cancer})}^{0.01}}{P(\text{Positive})}$$

Applying Bayes' Rule

- <u>Information:</u>
  - 1% of women in a given population have breast cancer
  - If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
  - If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

$$P(\text{Cancer} \mid \text{Positive}) = \frac{\overbrace{P(\text{Positive} \mid \text{Cancer})}^{0.9} \cdot \overbrace{P(\text{Cancer})}^{0.01}}{P(\text{Positive})}$$

Applying Bayes' Rule

- Information:
  - 1% of women in a given population have breast cancer
  - If a woman has breast cancer, there is a 90% chance that a particular diagnostic test will return a positive result (10% false negative rate)
  - If a woman does not have breast cancer, there is a 10% chance that this diagnostic test will return a positive result (10% false positive rate).

P(Positive) = P(True Positive) + P(False Positive)

P(Positive) = P(Positive | Cancer)· P(Cancer) + P(+ | Healthy)· P(Healthy)

P(Positive) = 0.9·0.01 + 0.1·(1-0.01)

P(Positive) = 0.108

Now we can complete Bayes' Rule

$$P(\text{Cancer} \mid \text{Positive}) = \frac{P(\text{Positive} \mid \text{Cancer}) \cdot P(\text{Cancer})}{P(\text{Positive})}$$

$$P(\text{Cancer} \mid \text{Positive}) = \frac{0.9 \cdot 0.01}{0.108} = 0.083$$

Frequentist Coin Flip: 20 Flips; 13 Heads

Objective: Estimate the Coin's Bias with a 95% Confidence Interval

```
binom.test(13, 20)

##
##  Exact binomial test
##
## data:  13 and 20
## number of successes = 13, number of trials = 20, p-value =
## 0.2632
## alternative hypothesis: true probability of success is not equal to 0.5
## 95 percent confidence interval:
##  0.4078115 0.8460908
## sample estimates:
## probability of success
##                   0.65
```

Conclusion:

- Bias = 0.65
- 95% CI = (0.41, 0.85)

Bayesian Coin Flip: 20 Flips; 13 Heads

Objective: Identify the bias (x) that yields the highest posterior probability. Given 13 heads were observed out of 20 flips

Likelihood

Prior Probability

Posterior Probability

$$P(\text{bias} = x \mid 13 \text{ heads}) = \frac{P(13 \text{ heads} \mid \text{bias} = x) \cdot P(\text{bias} = x)}{P(13 \text{ heads})}$$

Evidence or Marginal Likelihood

Bayesian Coin Flip: Define Priors

```
prior.probability <- numeric(101)
prior.probability[0:101] <- 1
# Normalize; since it is a PDF, sum must be 1.0
prior.probability <- prior.probability / (sum(prior.probability))
barplot(prior.probability, names.arg = coin.bias,xlab = "Coin Bias (x)", ylab = "P(bias = x)",ylim = c(0,0.02), m
ain = "Prior Probability Density Funciton: Biases Equally Likely", col ="#85C0F9")
```

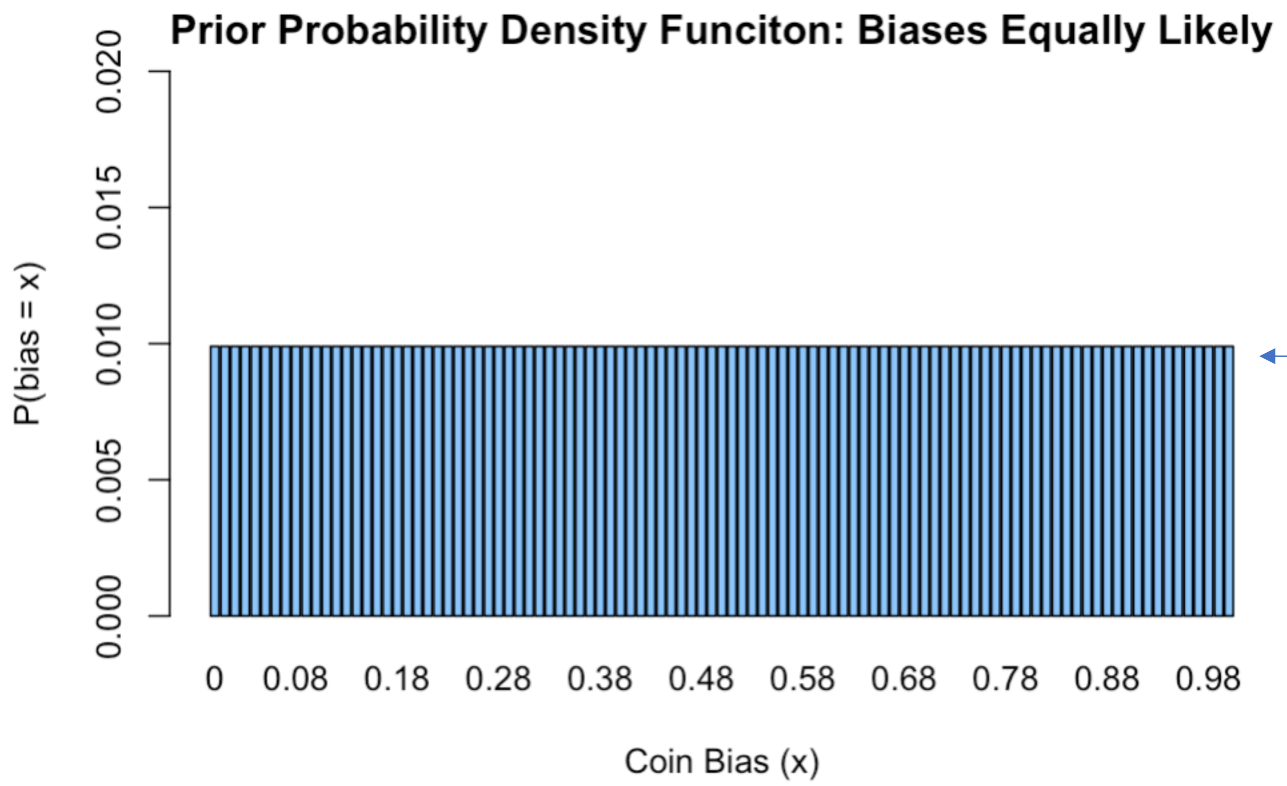

Prior for each potential bias
= 1/101
≅ 0.0099

Bayesian Coin Flip: Likelihood

```
coin.bias <- seq(from = 0, to = 1, by = 0.01)
likelihood <- dbinom(13, 20, prob = coin.bias)
barplot(likelihood, names.arg = coin.bias, ylab = "Likelihood: P(13 Heads | bias = x)", xlab = "Coin Bias (x)", c
ol = "#A95AA1") #Color-blindess friendly purple
```

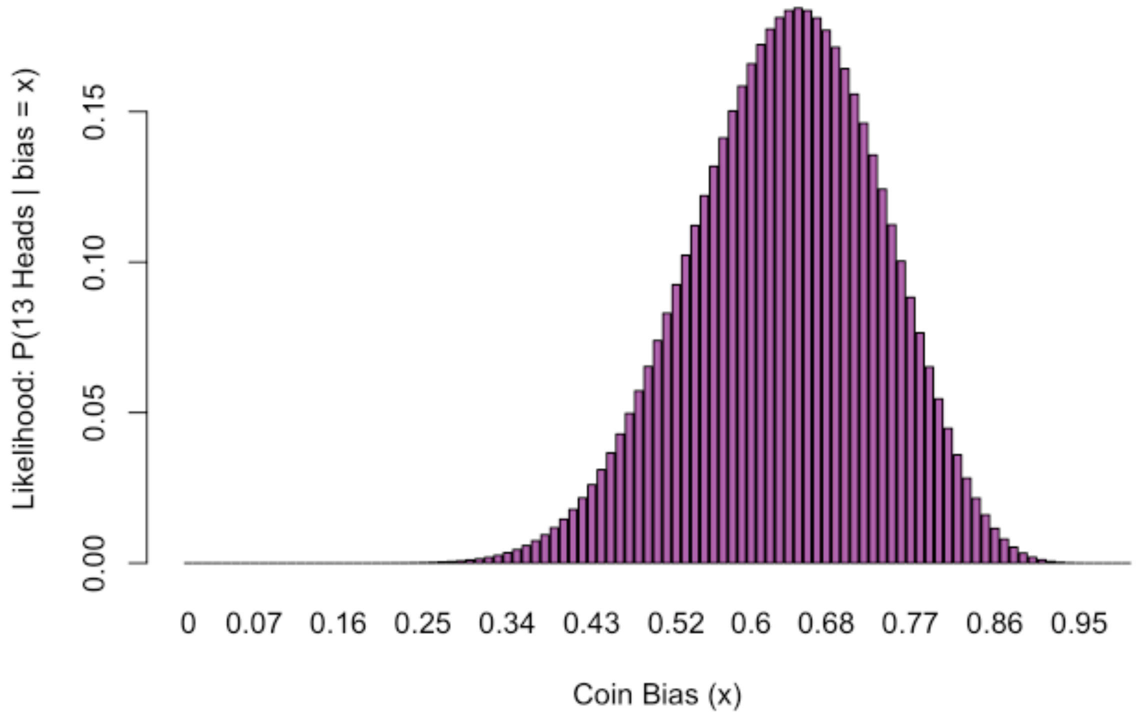

Marginal Likelihood

P(13 heads)

= P(13 heads | bias = 0.00) · P(bias = 0.00)

+ P(13 heads | bias = 0.01) · P(bias = 0.01)

+ P(13 heads | bias = 0.02) · P(bias = 0.02)

…

+ P(13 heads | bias = 0.50) · P(bias = 0.50)

…

+ P(13 heads | bias = 0.99) · P(bias = 0.99)

+ P(13 heads | bias = 1.00) · P(bias = 1.00)

Marginal Likelihood

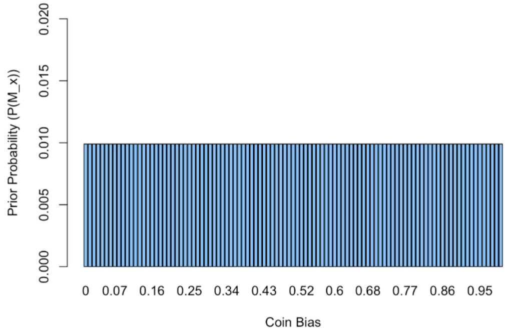

P(13 heads)

= P(13 heads | bias = 0.00) · P(bias = 0.00)

+ P(13 heads | bias = 0.01) · P(bias = 0.01)

+ P(13 heads | bias = 0.02) · P(bias = 0.02)

…

+ P(13 heads | bias = 0.50) · P(bias = 0.50)

…

+ P(13 heads | bias = 0.99) · P(bias = 0.99)

+ P(13 heads | bias = 1.00) · P(bias = 1.00)

Marginal Likelihood

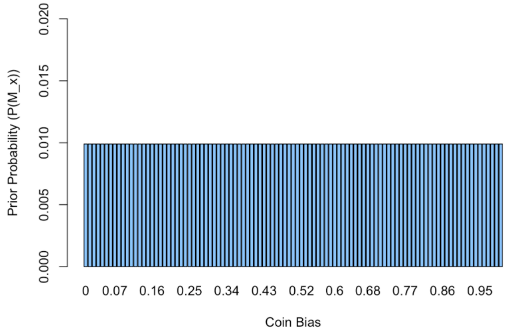

P(13 heads)

= P(13 heads | bias = 0.00) · 0.0099

+ P(13 heads | bias = 0.01) · 0.0099

+ P(13 heads | bias = 0.02) · 0.0099

…

+ P(13 heads | bias = 0.50) · 0.0099

…

+ P(13 heads | bias = 0.99) · 0.0099

+ P(13 heads | bias = 1.00) · 0.0099

Marginal Likelihood

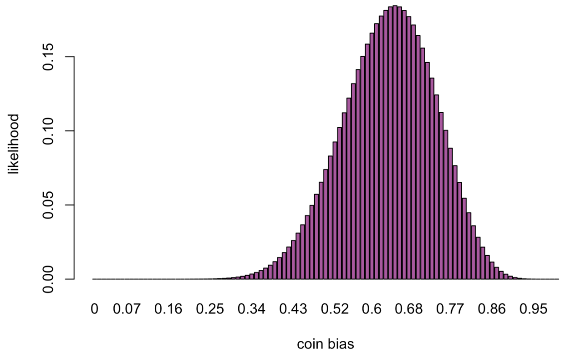

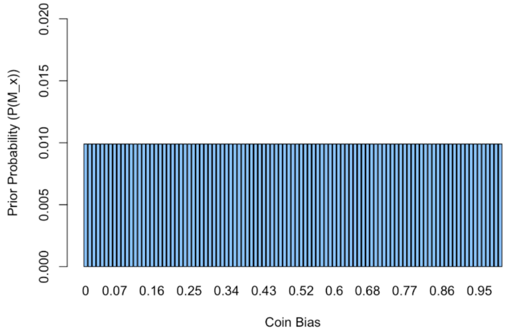

P(13 heads)

$= P(13 \text{ heads} \mid \text{bias} = 0.00) \cdot 0.0099$

$+ P(13 \text{ heads} \mid \text{bias} = 0.01) \cdot 0.0099$

$+ P(13 \text{ heads} \mid \text{bias} = 0.02) \cdot 0.0099$

…

$+ P(13 \text{ heads} \mid \text{bias} = 0.50) \cdot 0.0099$

…

$+ P(13 \text{ heads} \mid \text{bias} = 0.99) \cdot 0.0099$

$+ P(13 \text{ heads} \mid \text{bias} = 1.00) \cdot 0.0099$

Marginal Likelihood

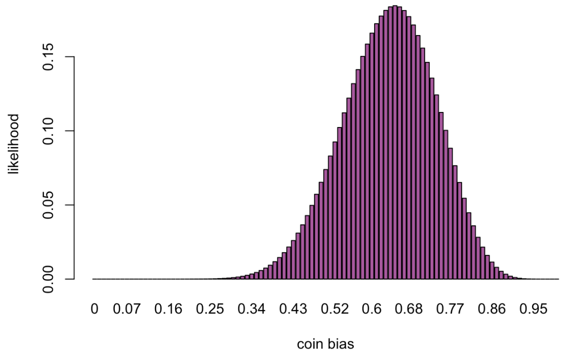

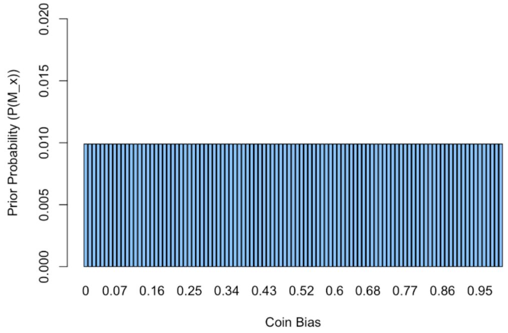

P(13 heads)

$= 0.0 \cdot 0.0099$

$+ 7.2e\text{-}22 \cdot 0.0099$

$+ 5.5e\text{-}18 \cdot 0.0099$

…

$+ 0.07392883 \cdot 0.0099$

…

$+ 6.8e\text{-}10 \cdot 0.0099$

$+ 0.0 \cdot 0.0099$

$= 0.04714757$

Marginal Likelihood

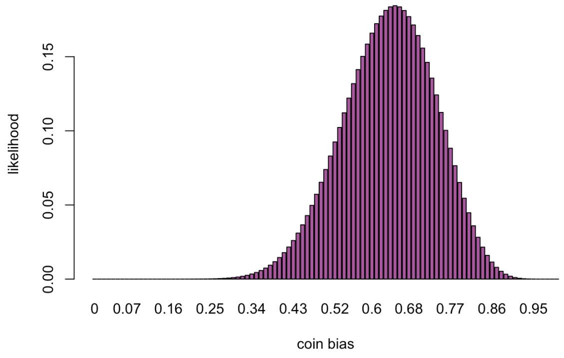

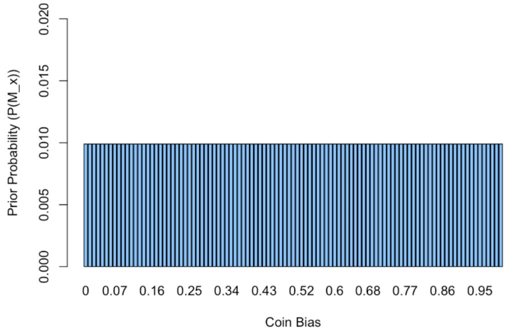


```
(p.d13 <- sum(dbinom(13, 20, coin.bias) * (1 / 101)))
```

```
## [1] 0.04714757
```

= 0.04714757

Posterior Probability

```
posterior.probability <- dbinom(13, 20, coin.bias) * (1 / 101) / p.d13
sum(posterior.probability)
```

```
## [1] 1
```

```
barplot(posterior.probability, names.arg = coin.bias, xlab = "Coin Bias (x)", y
lab = "Posterior Probability: P(bias = x | 13 Heads)", main = "Posterior Probab
ility Density Function: 13/20 Heads Observed")
```

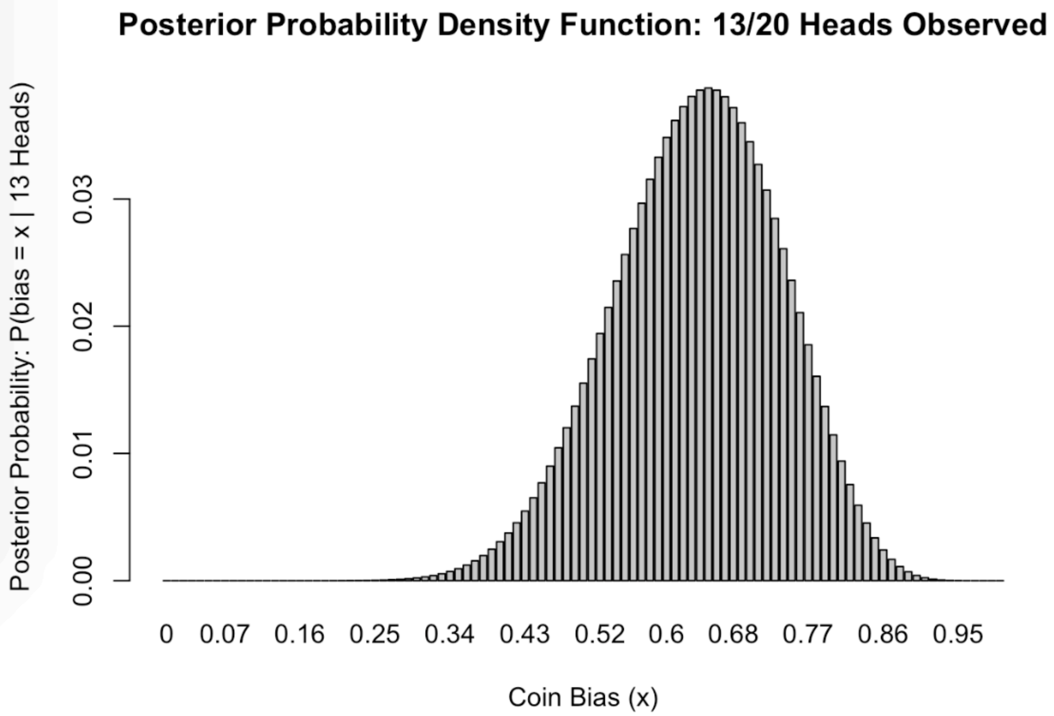

Recall Frequentist Conclusion:

- Bias = 0.65
- 95% CI = (0.41, 0.85)

Summary: Flipping a Coin with No expectations of fairness

$$P(\text{bias} = x \mid 13 \text{ heads}) = \frac{P(13 \text{ heads} \mid \text{bias} = x) \cdot P(\text{bias} = x)}{P(13 \text{ heads})}$$

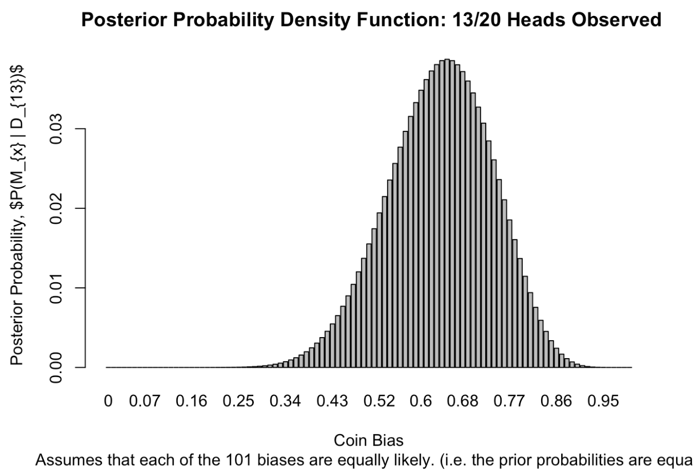

=

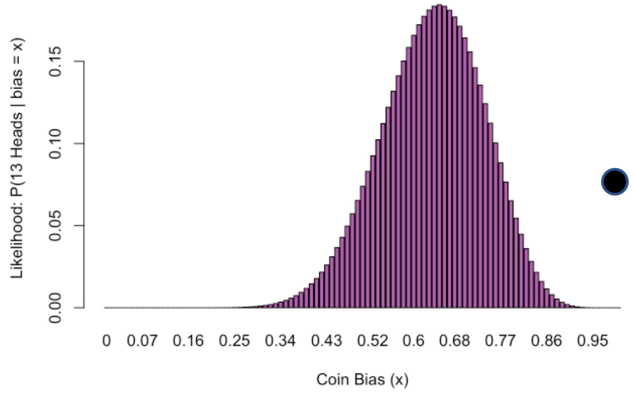

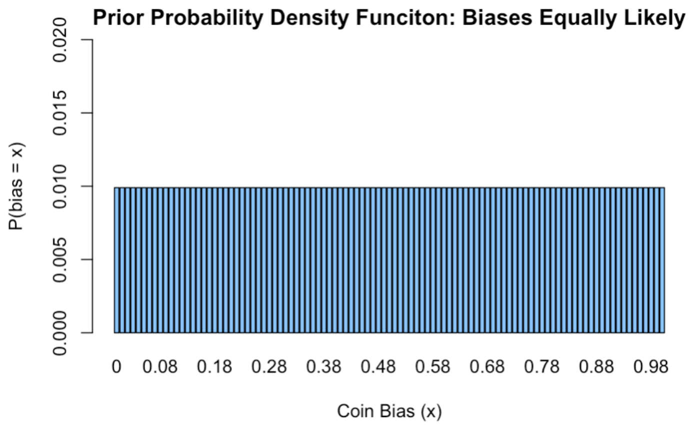

P(13 Heads) = 0.04714757

Going Further