Bayes theory

The Hypotheses

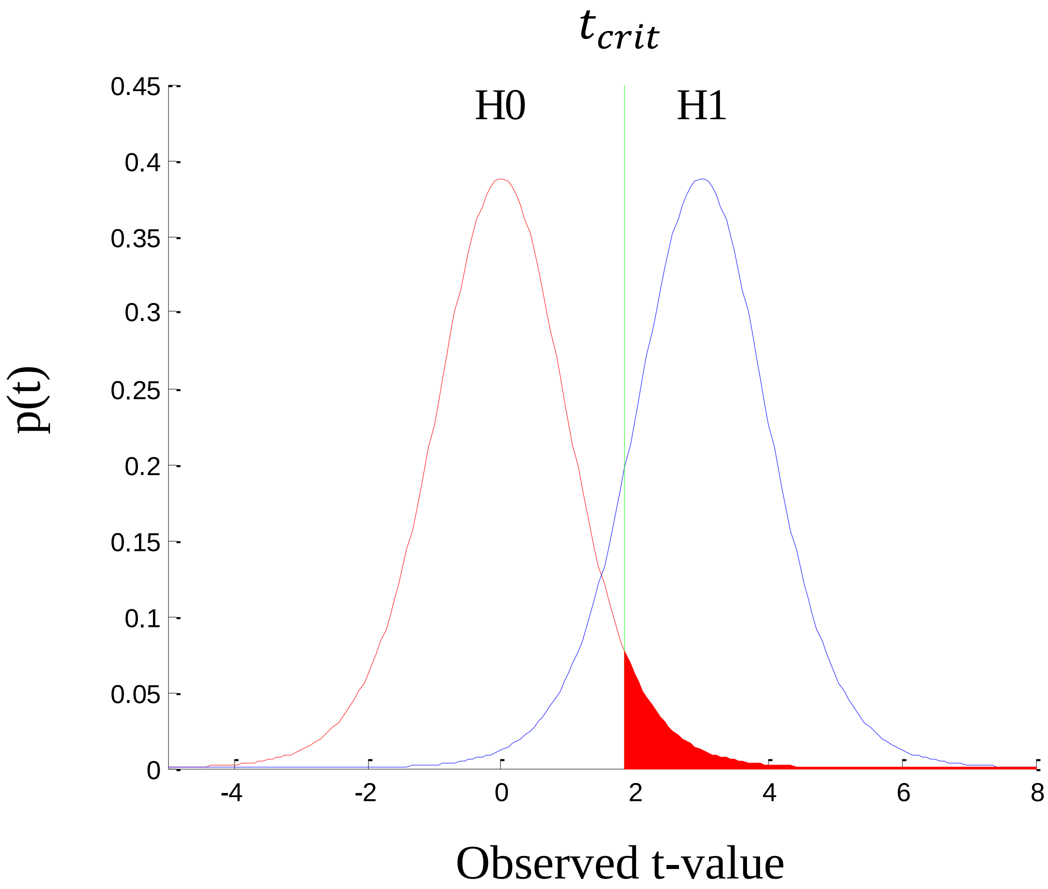

	H0 True	H1 True
We Say H0 “True”	1-α	β
We Say H1 True	α	Power (1-β)

The Hypotheses

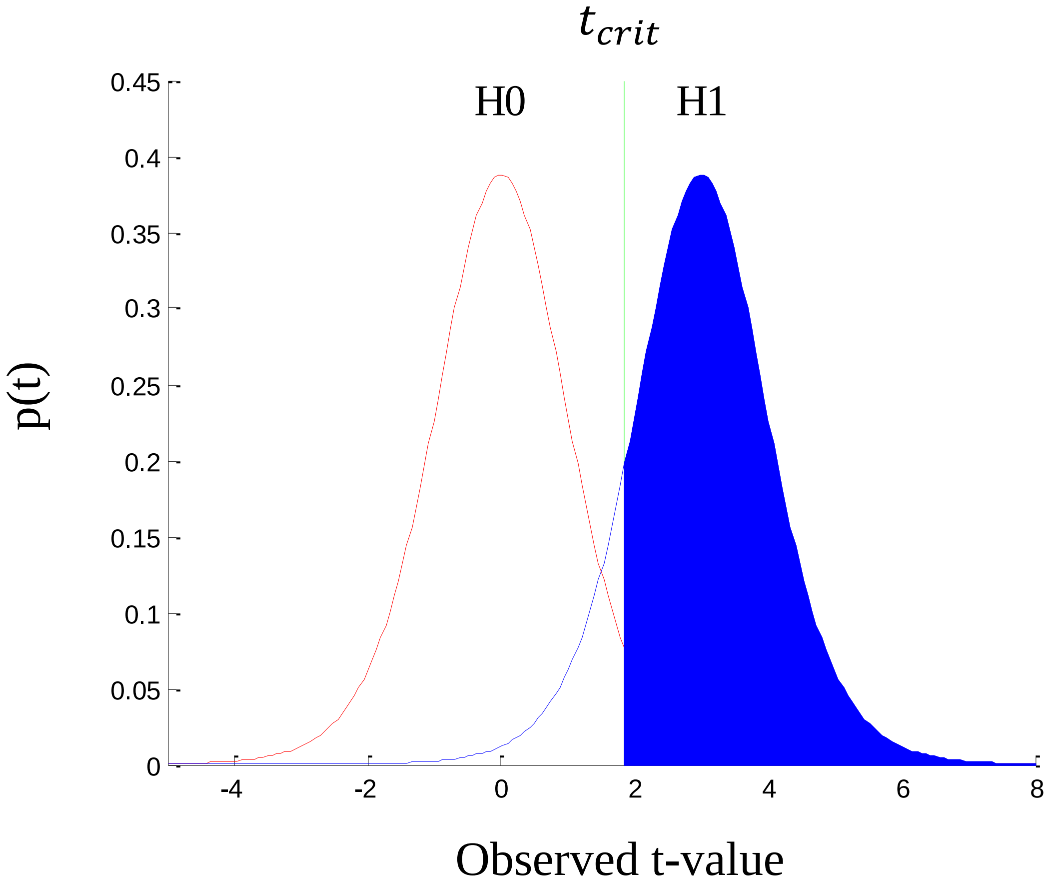

	H0 True	H1 True
We Say H0 “True”	1-α	β
We Say H1 True	α	Power (1-β)

Discriminability is independent of Criterion

Discriminability (d') = $|\mu_{SN} - \mu_N| / \sigma_N$

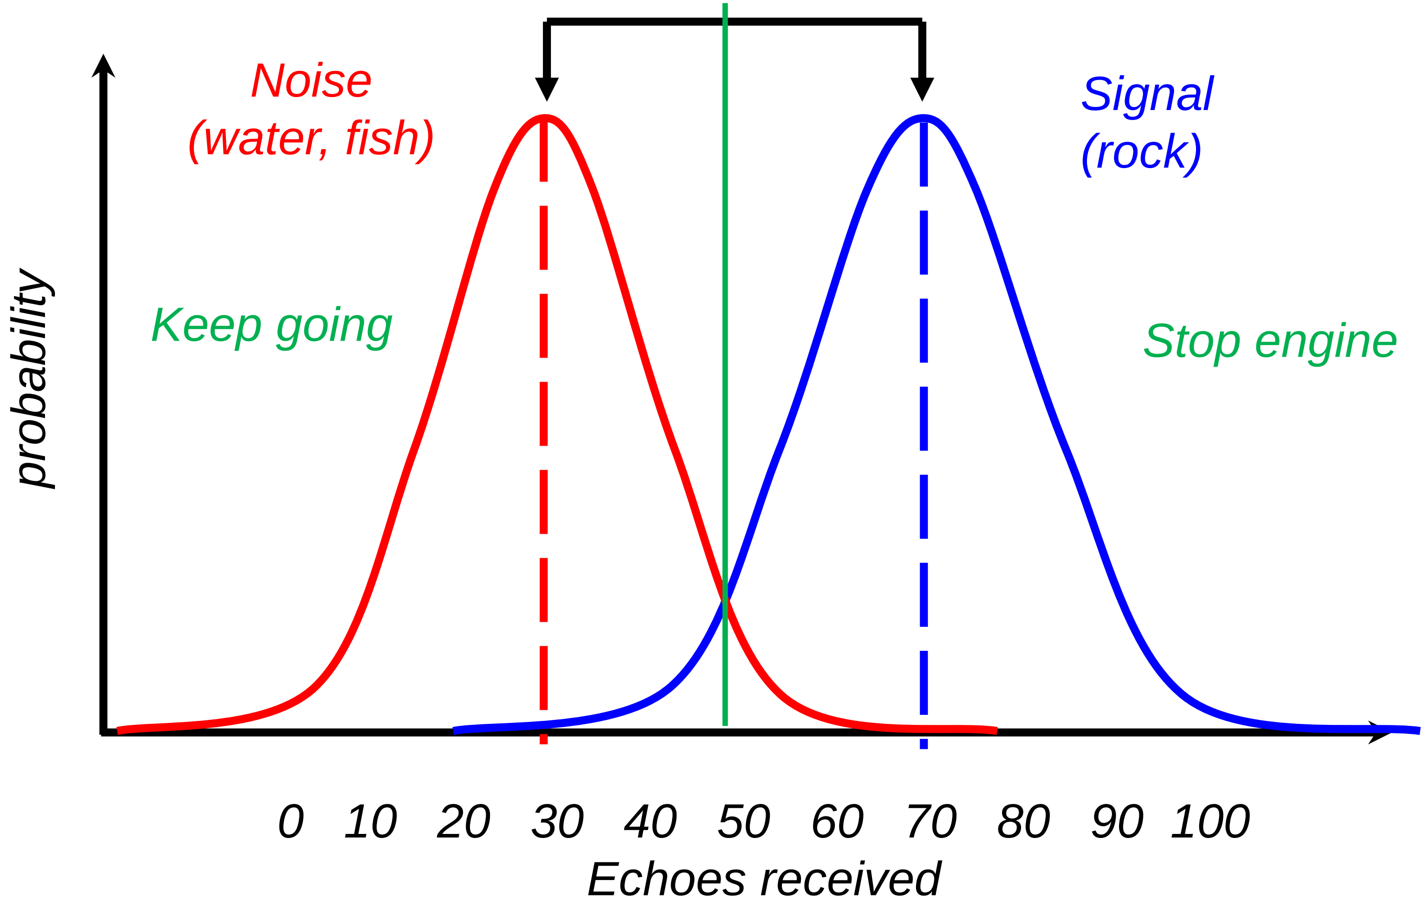

When Does NHST Make Optimal Decisions?

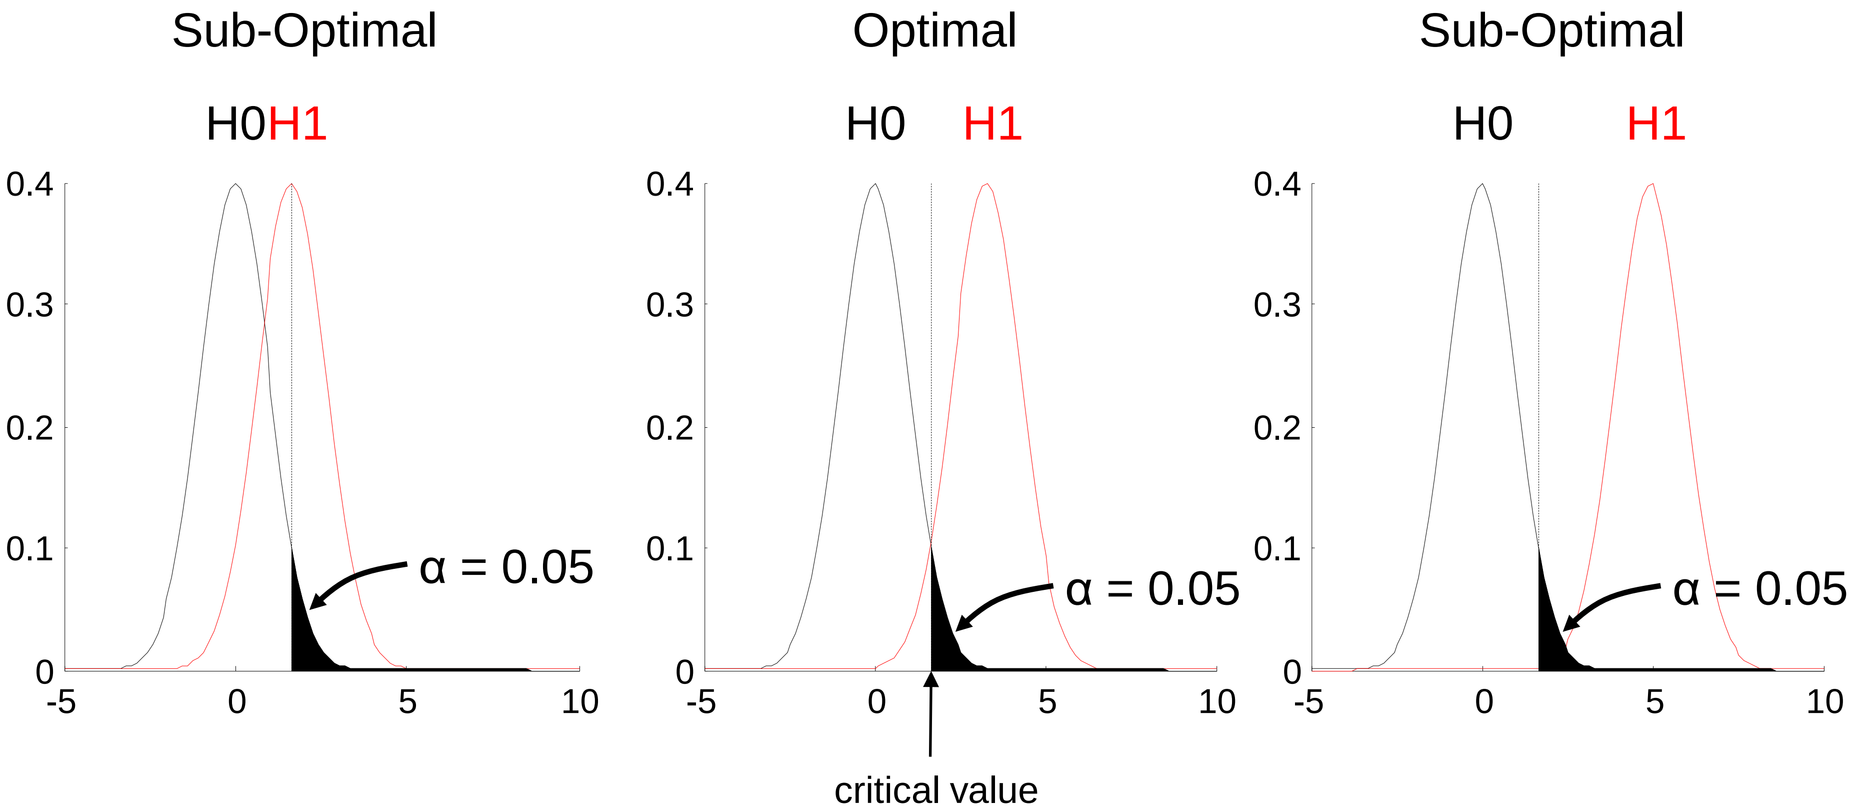

Null-hypothesis significance testing only makes optimal decisions when the intersection between the H0 and the H1 distributions overlaps with the critical value.

Bayes Theorem

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

$$\text{Posterior} = \frac{\text{Likelihood} \times \text{Prior}}{\text{Normalizer}}$$

The Problem

- Suppose you have isolated the DNA from a prehistoric tree and managed to make clones in your lab.
- You grow a few trees and measure their dimensions.

The Problem

Population:

X
1
2
2
3
3
3
4
4
5

The Problem

- A friend calls you up and tells you they have also isolated the DNA from some prehistoric trees and made clones and they want to know if the trees are from the same species.

The Problem

x
5
6

Calculate the probability that our friends sample mean came from our distribution of sample means

$$t = \frac{m - \mu}{s/\sqrt{n}}$$

$$t = \frac{5.5 - 3}{0.71/\sqrt{2}}$$

$$t = 5$$

There is a 6% chance of observing a sample mean as extreme or more extreme than 5.5 meters from our tree height population (t(1) = 5, p = 0.06). Decision: there is not enough evidence to reject that H0 hypothesis that our friend's tree was the from the same population as our trees.

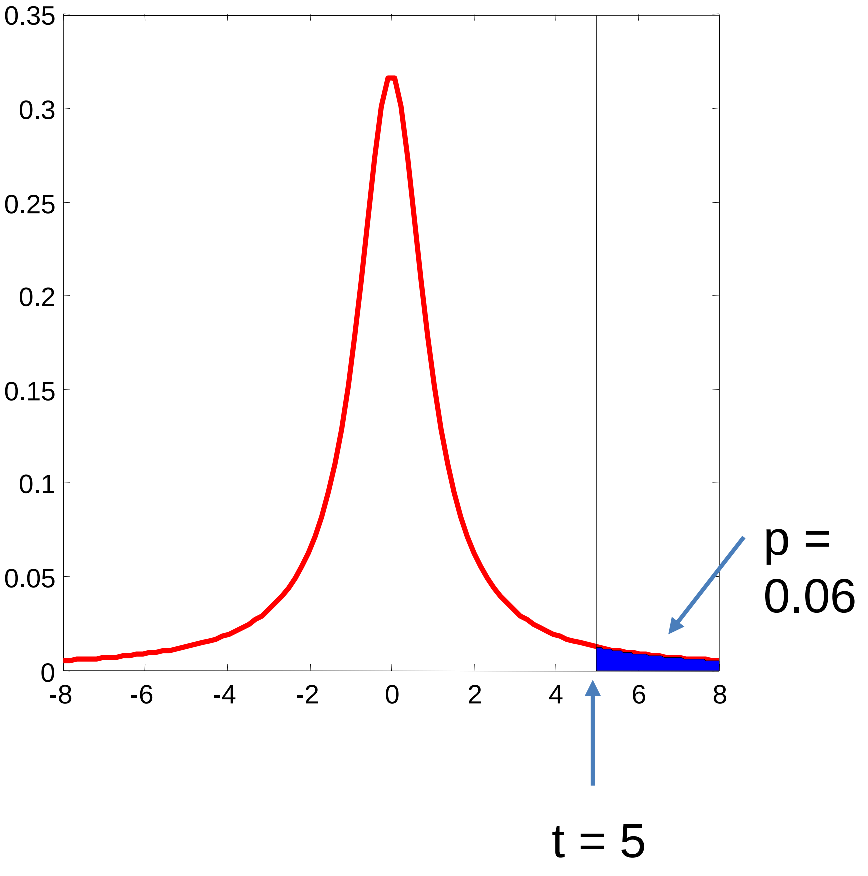

Can we say more about our friend’s sample if we assume we have made a statistically optimal decision?

- Yes!
- Using Bayes theorem we can calculate the probability that the null hypothesis (H_0) is true and the probability that the alternative hypothesis (H_1) is true and compare them.

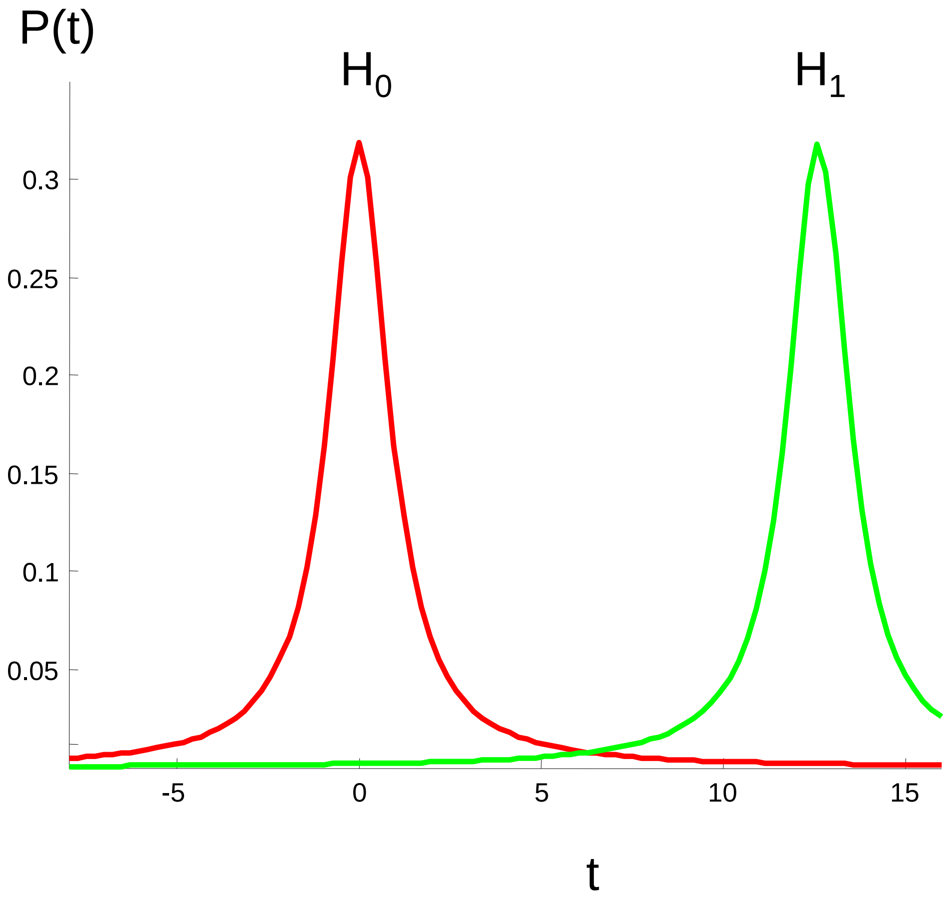

The Bayesian Approach

- $p(H_0|t) = \frac{p(t|H_0)p(H_0)}{p(t)}$
- $p(H_0|t=5) = \frac{0.0122 \cdot 0.5}{p(t)}$

- $p(H_1|t) = \frac{p(t|H_1)p(H_1)}{p(t)}$
- $p(H_1|t=5) = \frac{0.0054 \cdot 0.5}{p(t)}$

- $\frac{p(H_0|t)}{p(H_1|t)} = \frac{\frac{0.0122 \cdot \cancel{0.5}}{\cancel{p(t)}}}{\frac{0.0054 \cdot \cancel{0.5}}{\cancel{p(t)}}}$
- $\frac{p(H_0|t)}{p(H_1|t)} = \frac{0.0122}{0.0054}$=2.76

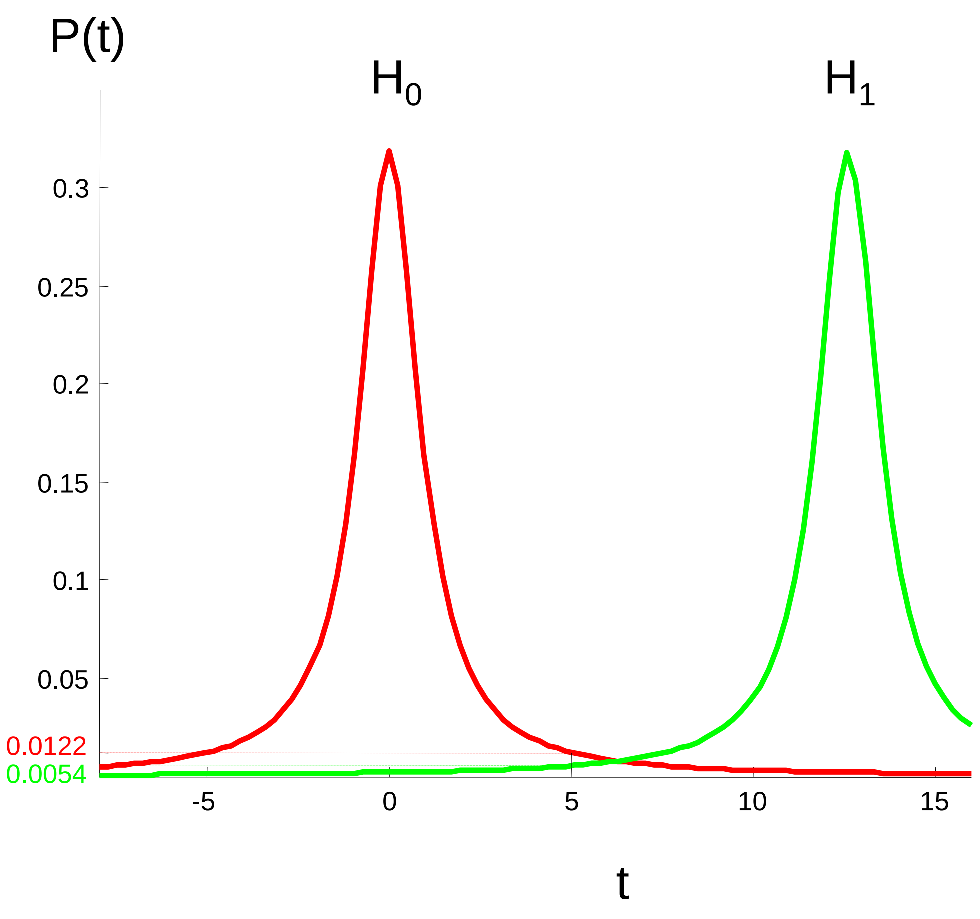

The Bayesian Approach

- The odds are 2.76:1 that our friends sample of trees came from the same population as our trees.

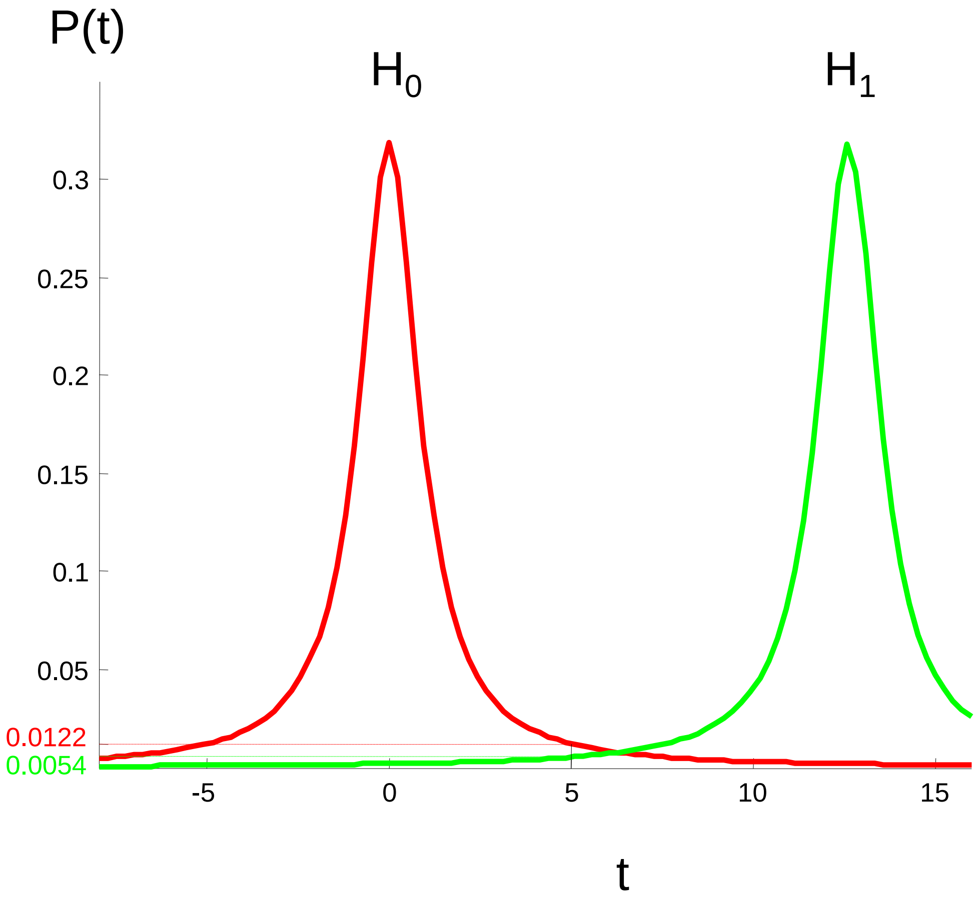

Comparison

- Null Hypothesis Significance Testing
  - Requires likelihood distributions.
  - Implicitly assumes prior distributions.
  - Assumes that the null and alternative hypothesis population distributions are Gaussian and have the same variance.
  - Makes statistically optimal decisions in very limited cases.
- Bayesian Approach
  - Requires likelihood distributions.
  - Requires prior distributions.
  - Always makes statistically optimal decisions.

Problem: there is 1 Null and infinitely many H1 hypotheses

Example 1

- A fortune teller claims to be able to see into the future.
- You devise an experiment where the fortune teller tries to predict whether stocks will rise or fall on a given day.
- You have the fortune teller make predictions about 100 stocks and 48 end up being correct.

Example 1: Null Hypothesis Significance Testing

- H0: The number of stocks the fortune teller can accurately predict is not significantly different from chance.
- H1: The number of stocks the fortune teller can accurately predict is significantly different from chance.
- We'll constrain our Type I error rate to $\alpha = 0.05$.

Example 1: Bayesian Inference

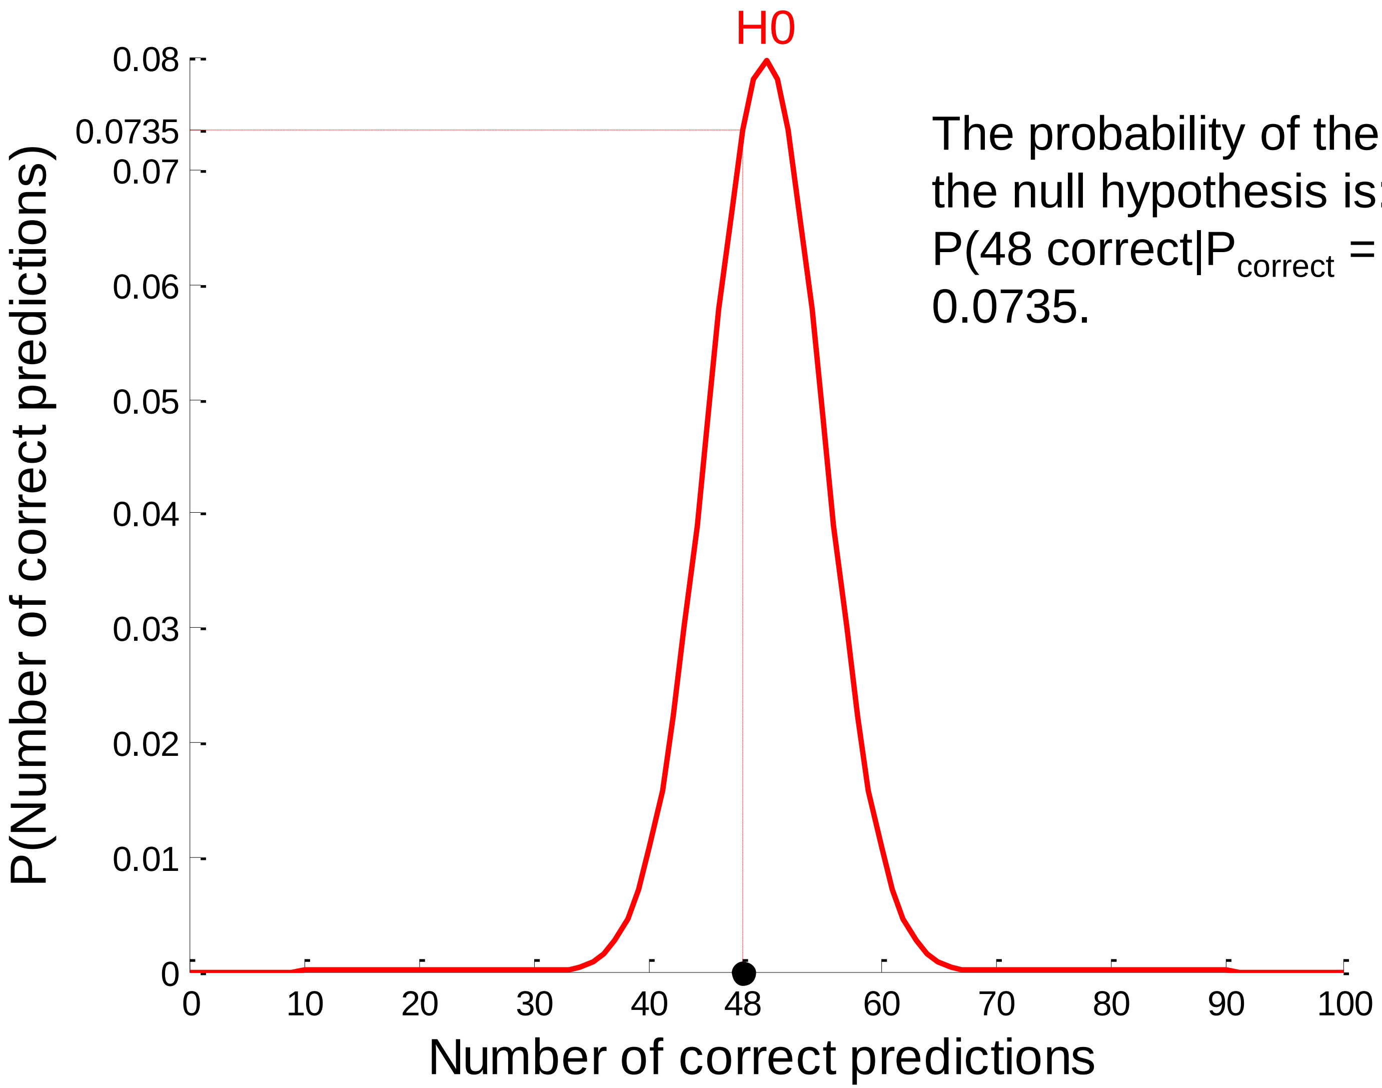

The probability of the data under the null hypothesis is: $P(48 \text{ correct}|P_{correct} = 0.5) = 0.0735$.

Example 1: Bayesian Inference

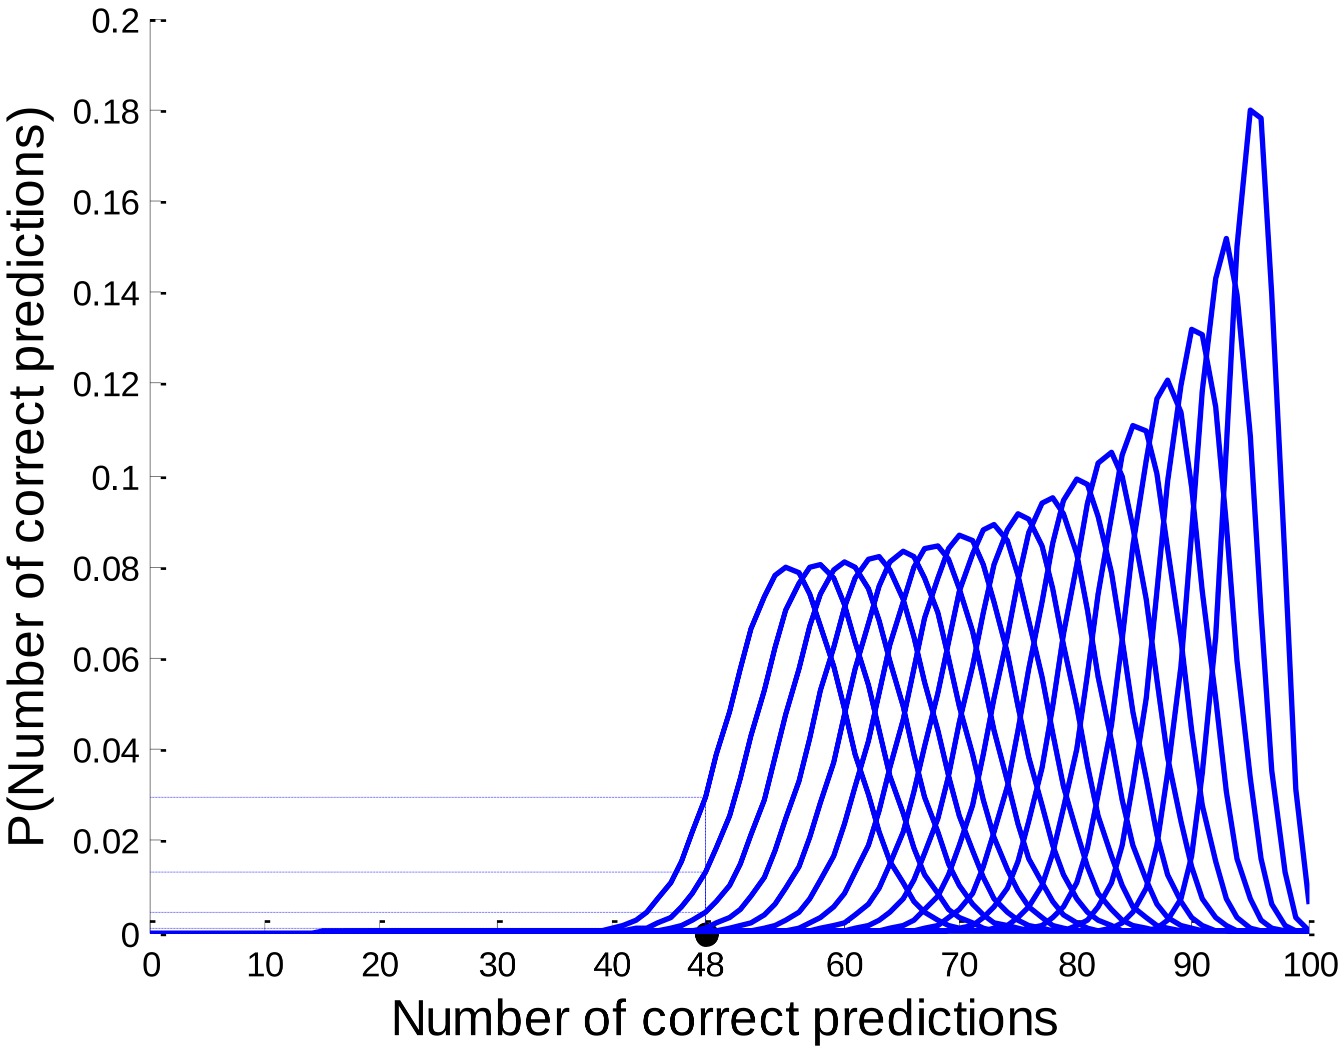

Example 1: NHST

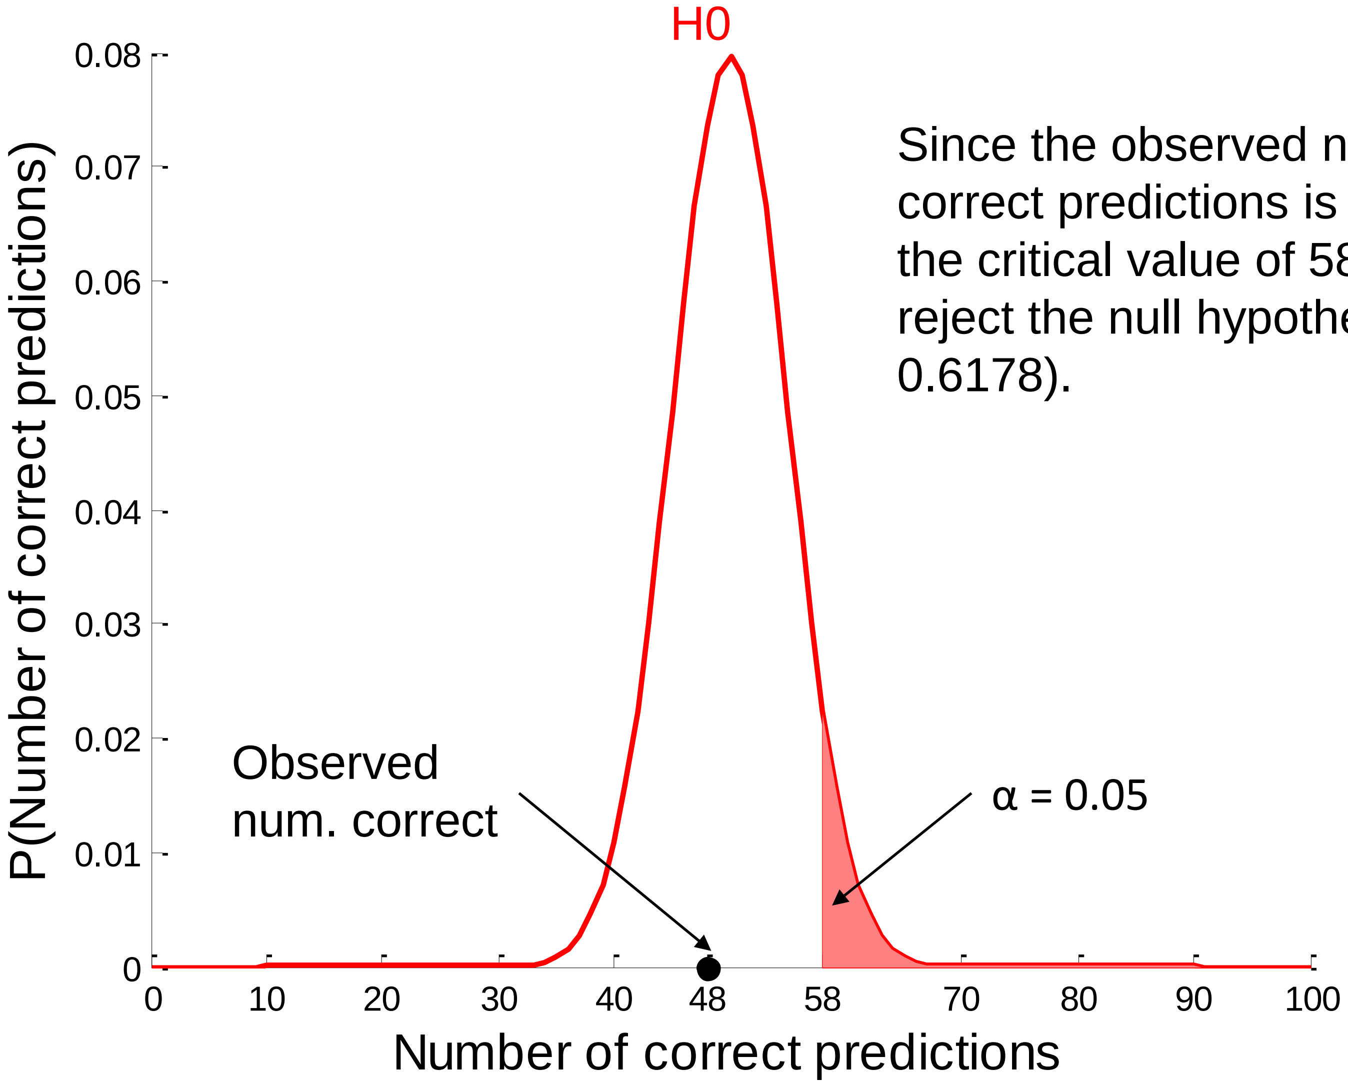

Since the observed number of correct predictions is less than the critical value of 58, we fail to reject the null hypothesis (p = 0.6178).

Example 1: Bayesian Inference

- H0: The fortune teller is at chance when predicting stock price changes.
  - $P_{correct} = 0.5$
- H1: The fortune teller makes predictions that are better than chance.
  - $P_{correct} > 0.5$

Example 1: Bayesian Inference

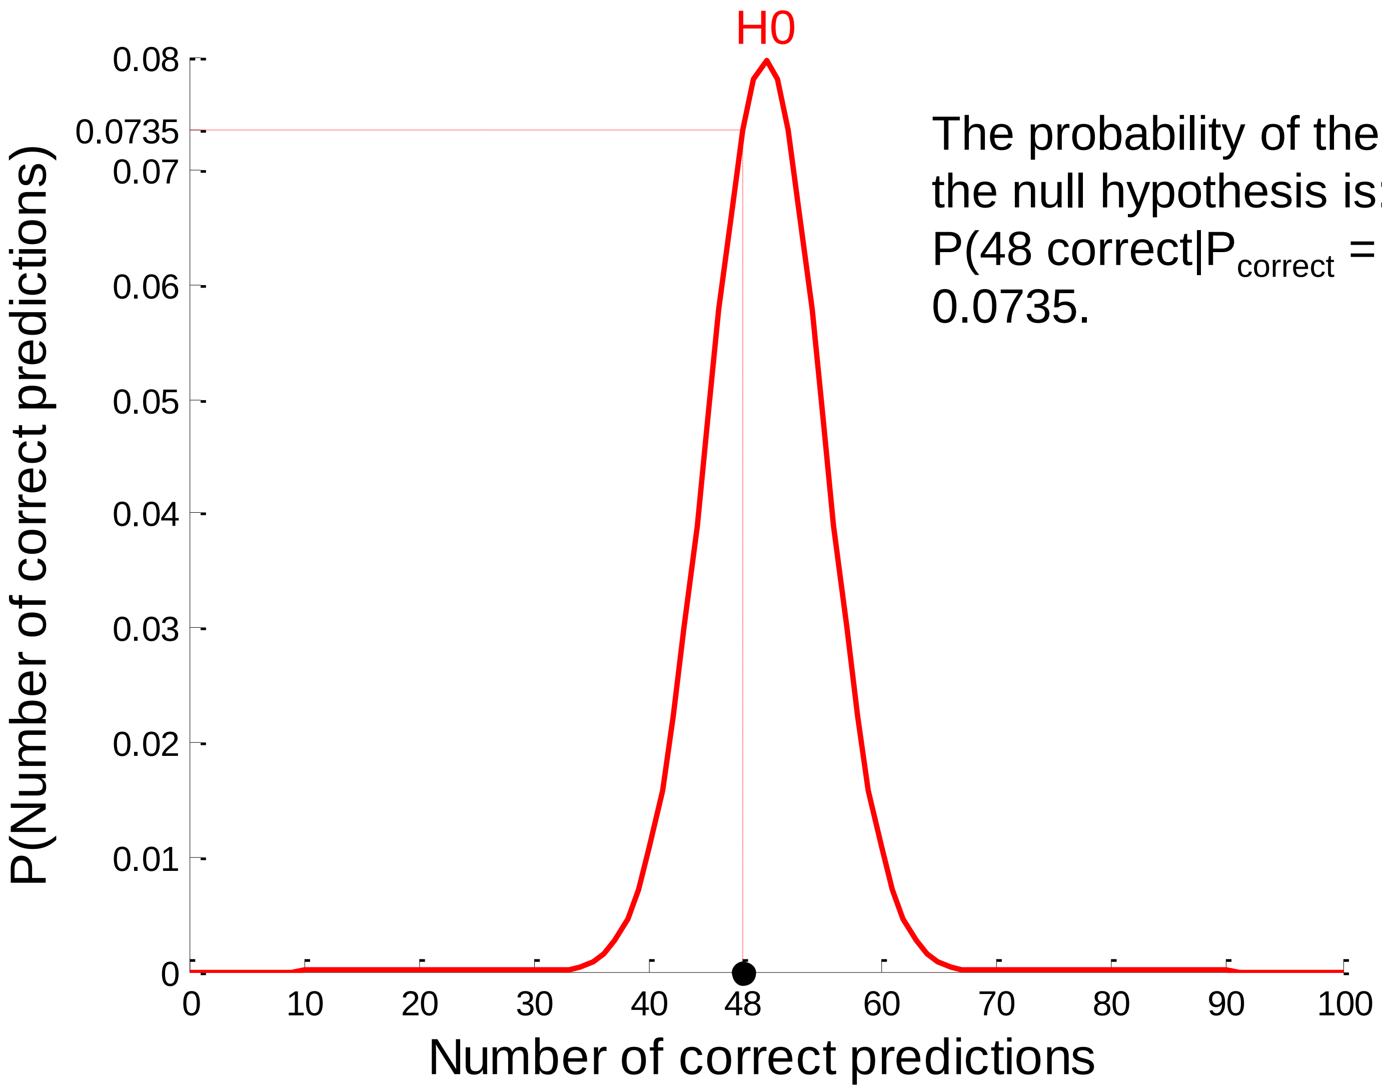

The probability of the data under the null hypothesis is: $P(48\ \text{correct}|P_{correct} = 0.5) = 0.0735$.

Example 1: Bayesian Inference

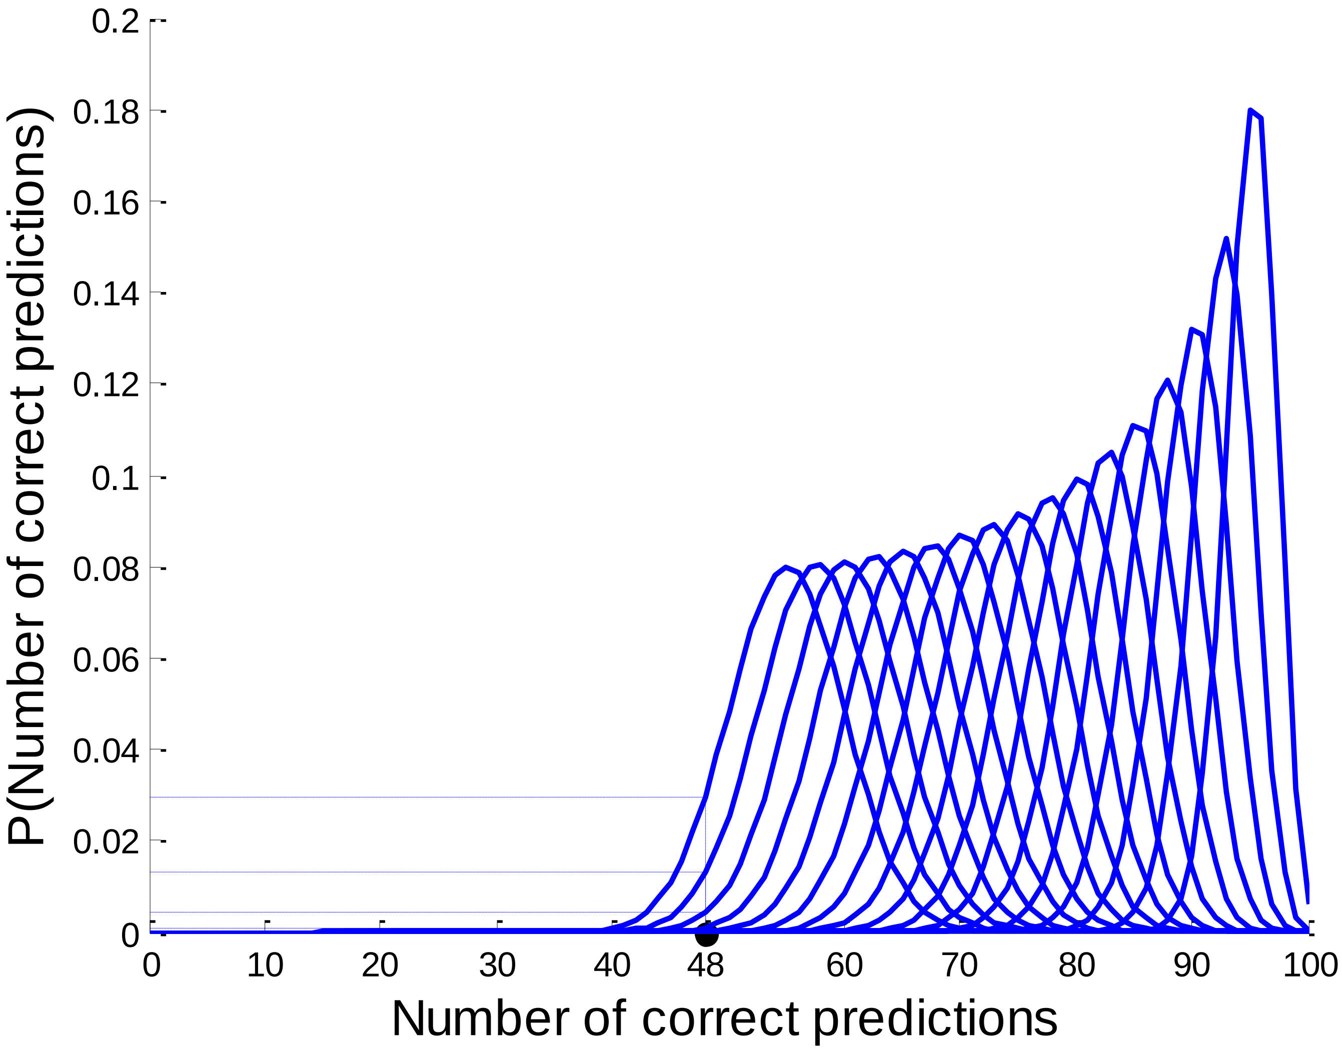

Example 1: Bayesian Inference

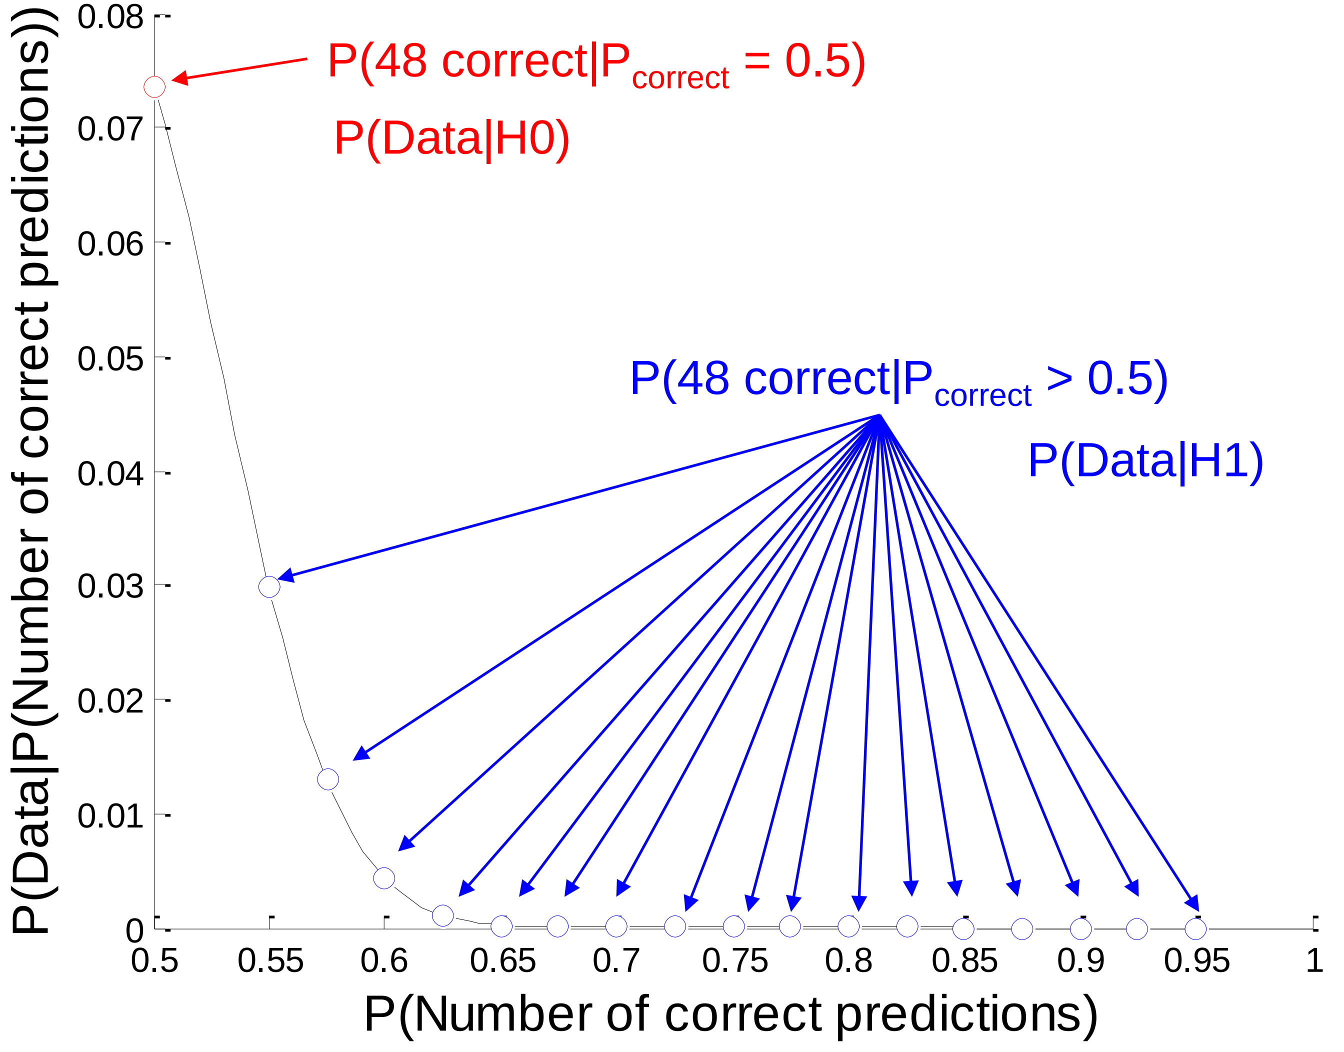

Example 1: Bayesian Inference

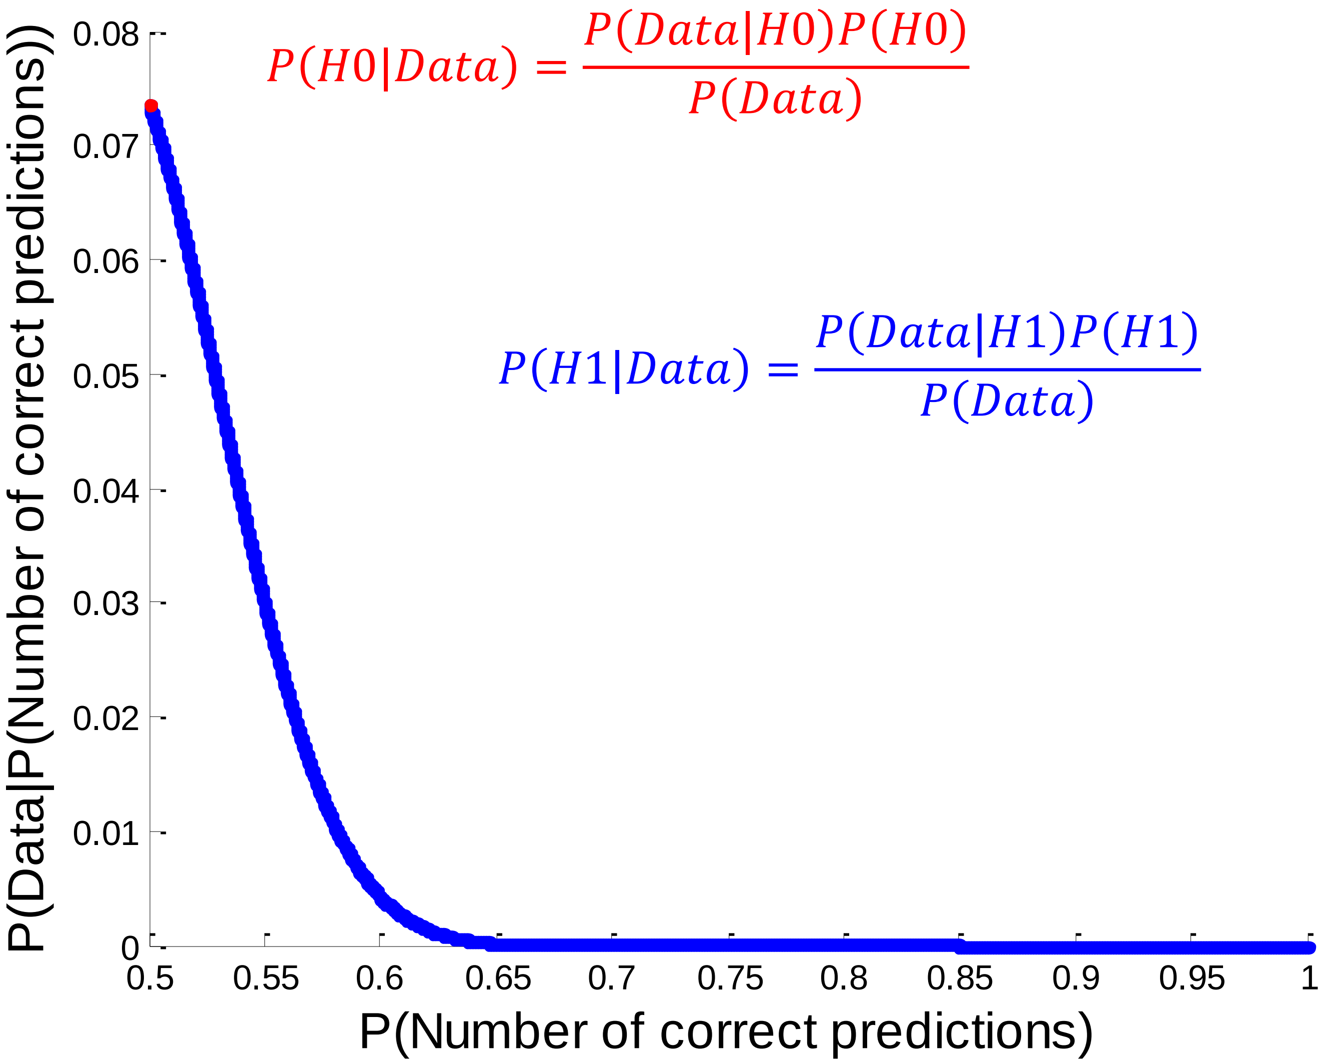

Example 1: Bayesian Inference

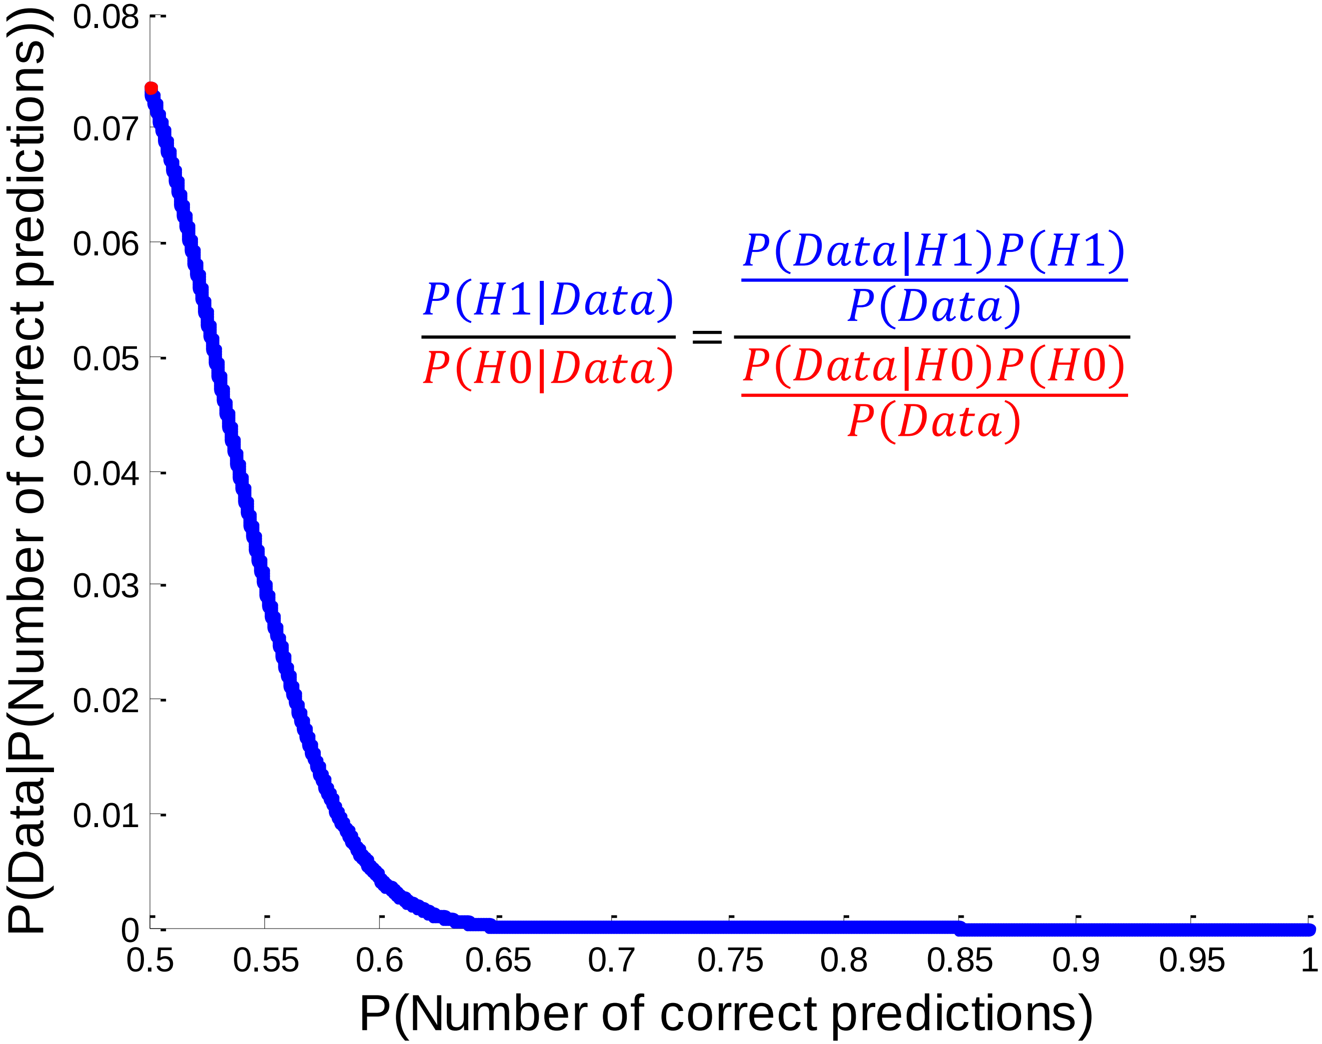

$$\frac{P(H1|Data)}{P(H0|Data)} = \frac{\frac{P(Data|H1)P(H1)}{P(Data)}}{\frac{P(Data|H0)P(H0)}{P(Data)}}$$

Example 1: Bayesian Inference

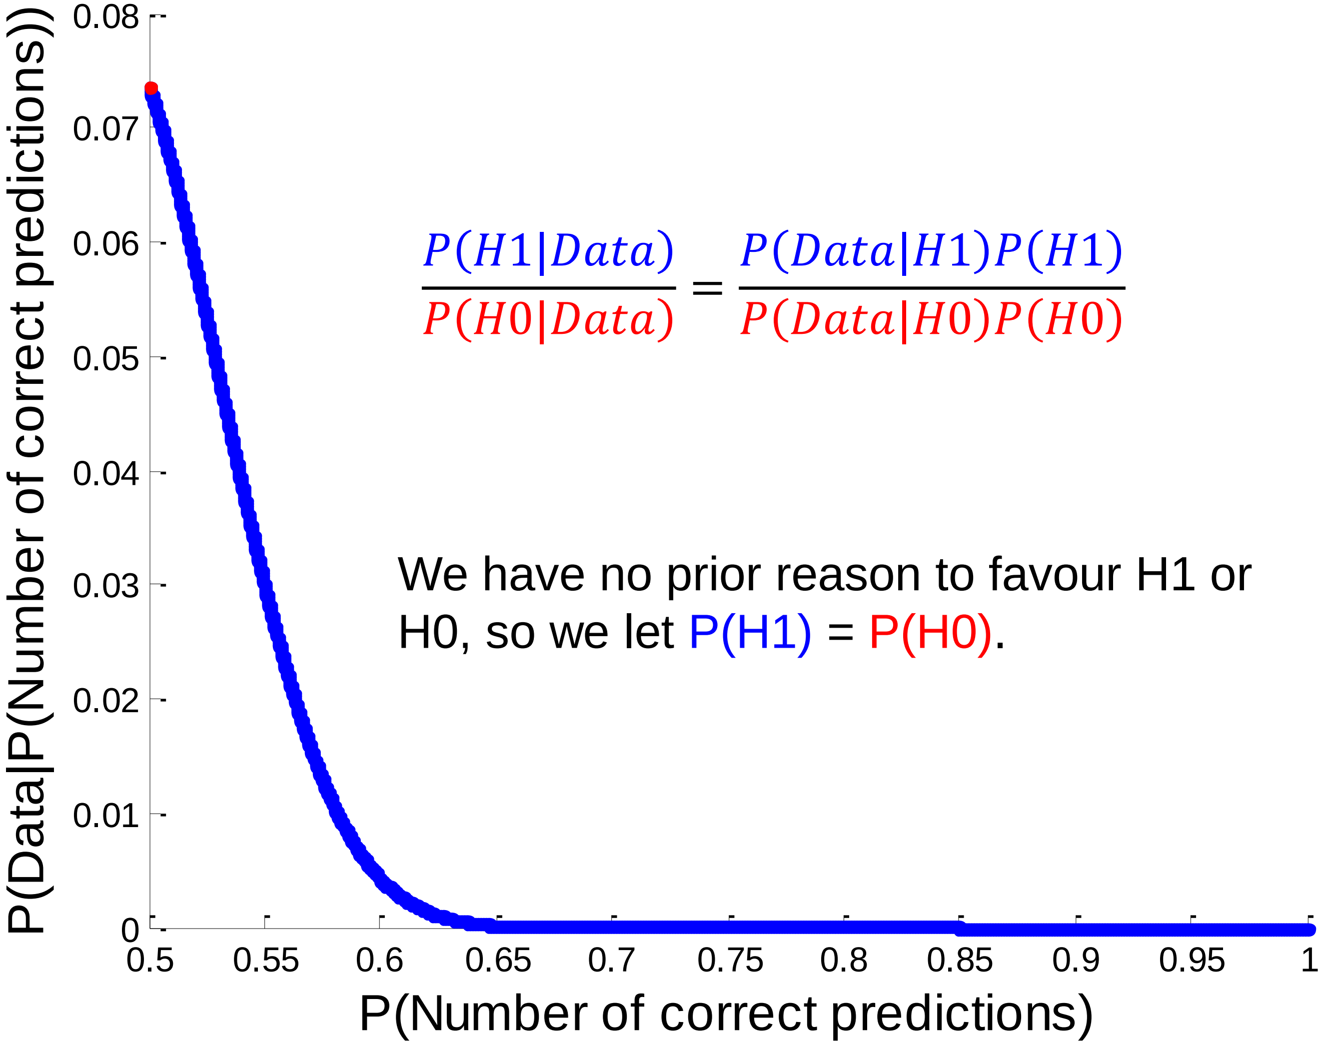

$$\frac{P(H1|Data)}{P(H0|Data)} = \frac{P(Data|H1)P(H1)}{P(Data|H0)P(H0)}$$

We have no prior reason to favour H1 or H0, so we let P(H1) = P(H0).

Example 1: Bayesian Inference

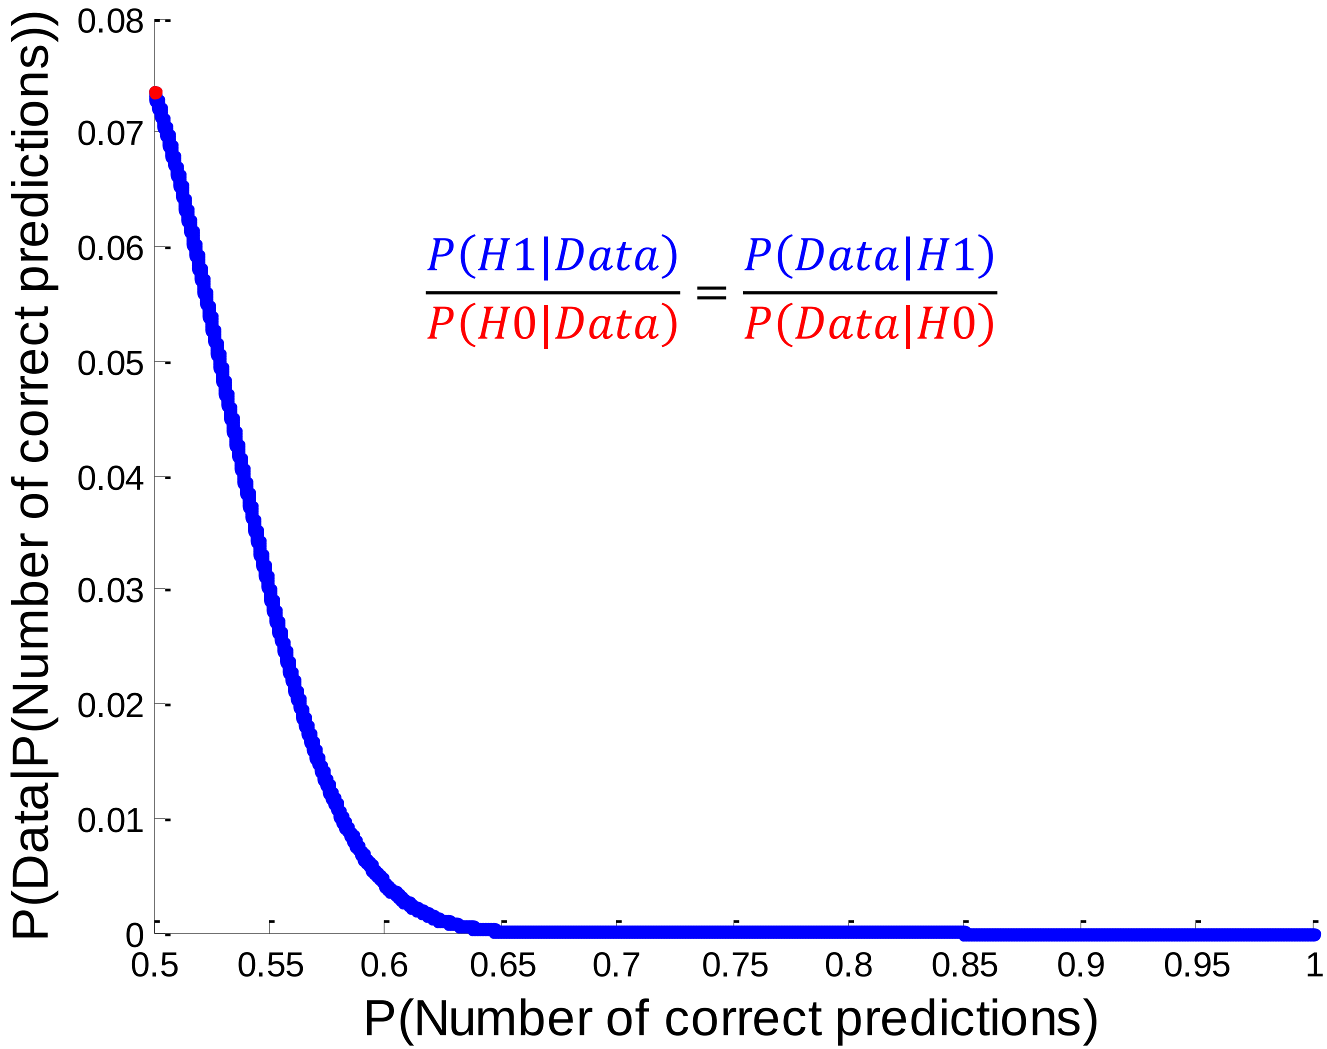

Example 1: Bayesian Inference

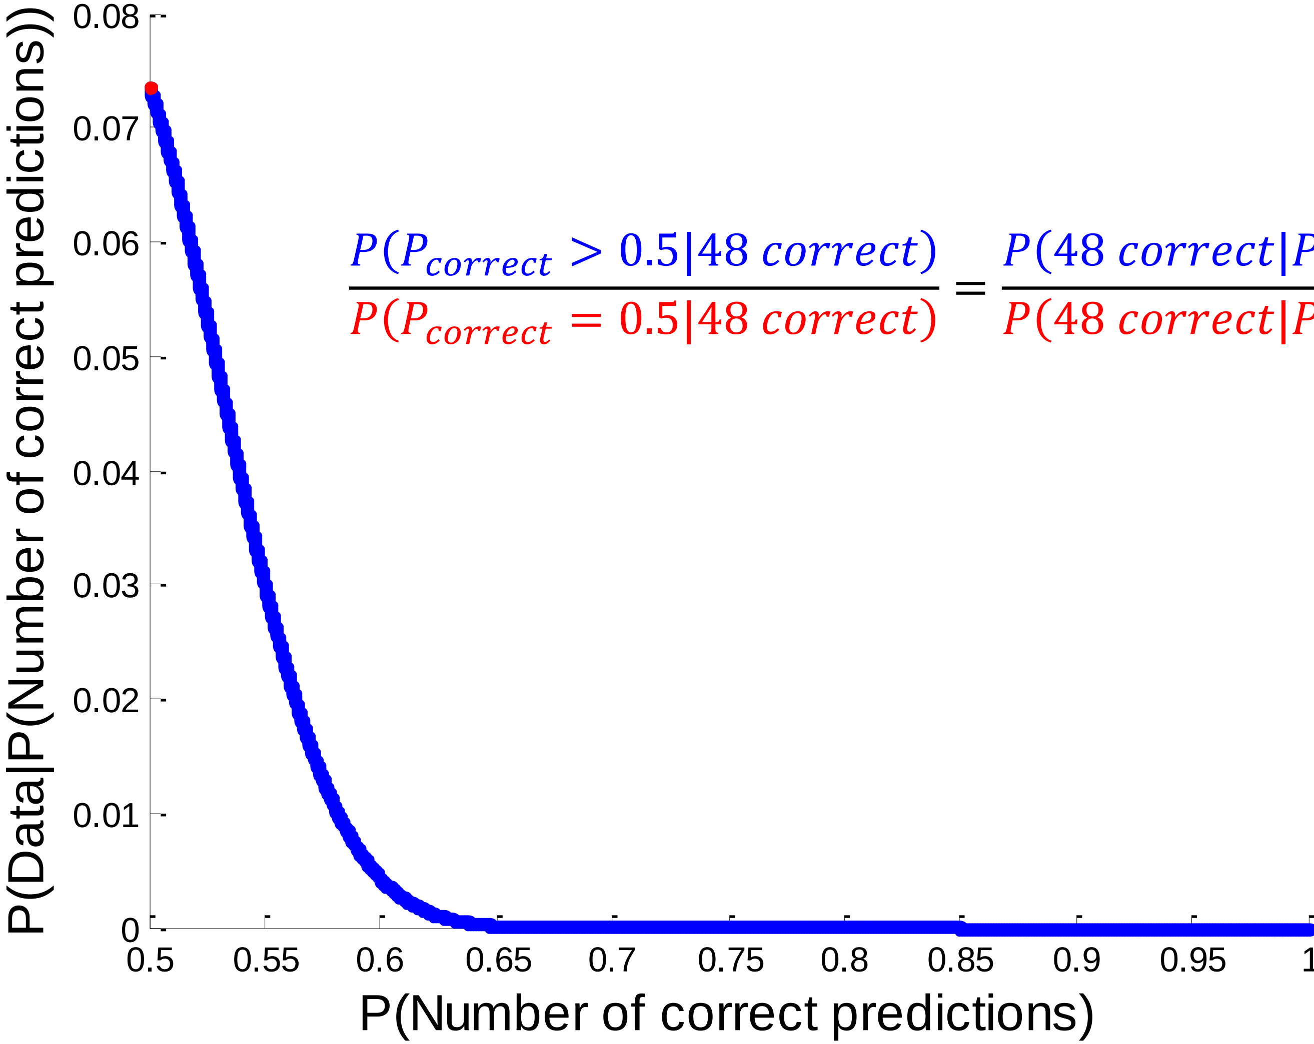

$$\frac{P(P_{correct} > 0.5|48\ correct)}{P(P_{correct} = 0.5|48\ correct)} = \frac{P(48\ correct|P_{correct} > 0.5)}{P(48\ correct|P_{correct} = 0.5)}$$

Example 1: Bayesian Inference

$$\frac{P(P_{correct} > 0.5|48\ correct)}{P(P_{correct} = 0.5|48\ correct)} = \frac{\int_{P_{correct}=0.5+d\tau}^{1} P(48\ correct|P_{correct} = \tau)d\tau}{0.0735}$$

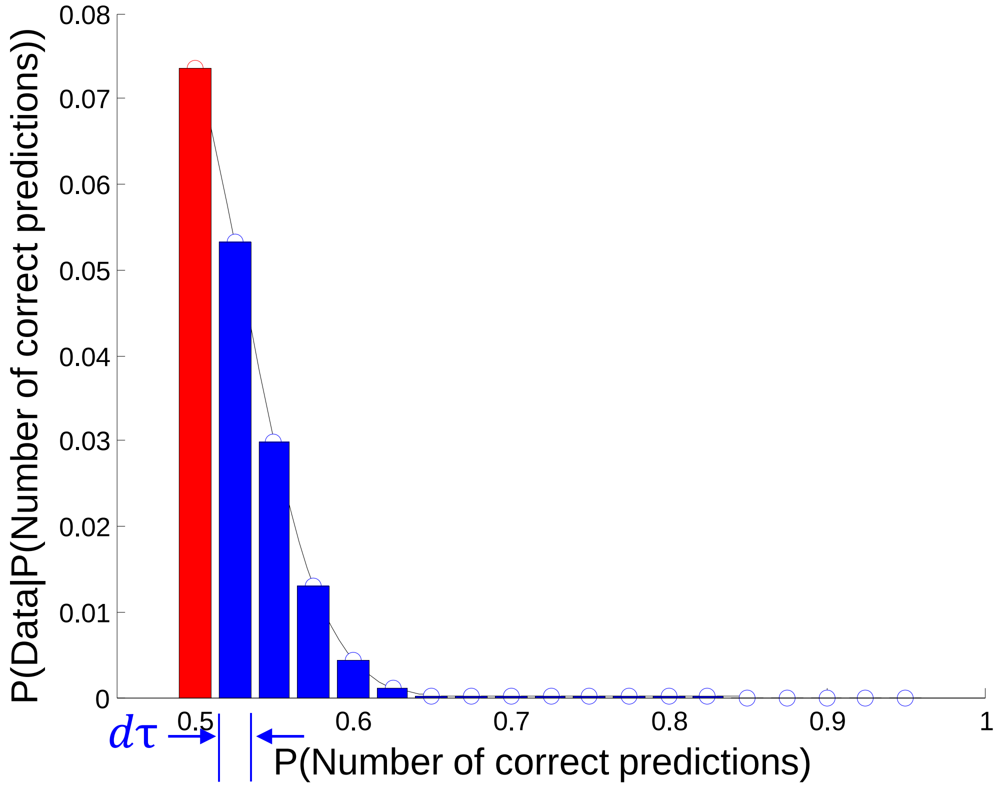

Example 1: Bayesian Inference

$$\frac{P(P_{correct} > 0.5|48\ correct)}{P(P_{correct} = 0.5|48\ correct)} = \frac{\left[\sum_{P_{correct}>0.5} P(48\ correct|P_{correct}) \cdot d\tau\right]/n_{bins}}{0.0735}$$

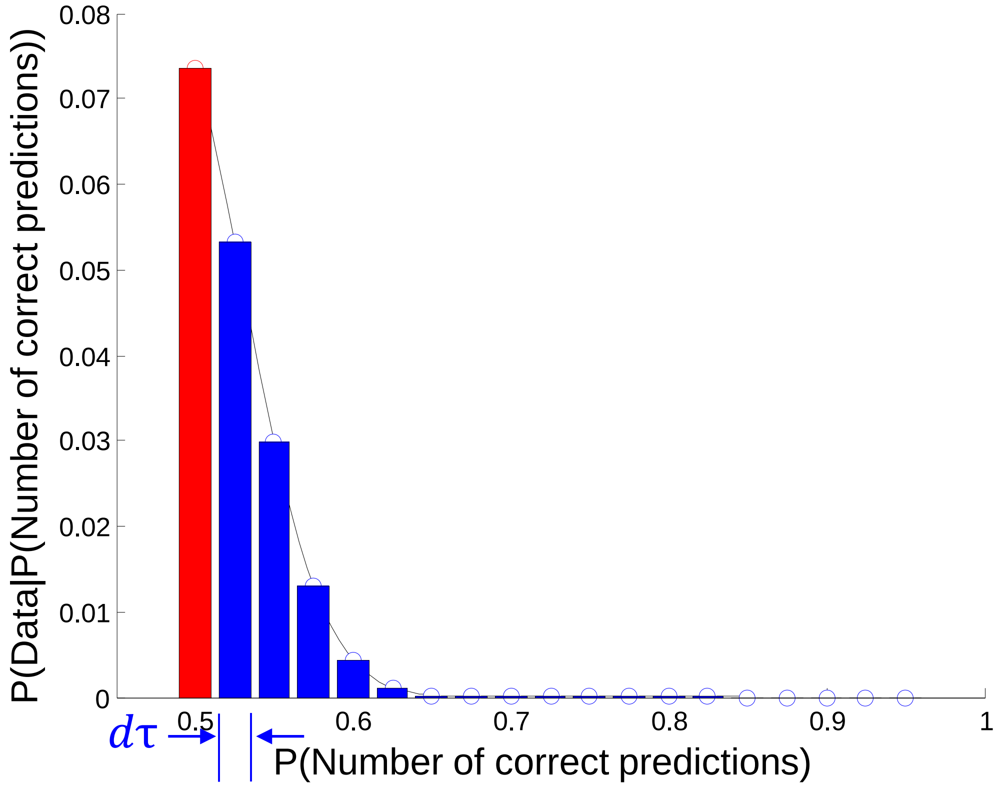

Example 1: Bayesian Inference

$$\frac{P(P_{correct} > 0.5|48\ correct)}{P(P_{correct} = 0.5|48\ correct)} = 0.0930$$

Using n_{bins} = 100 000.

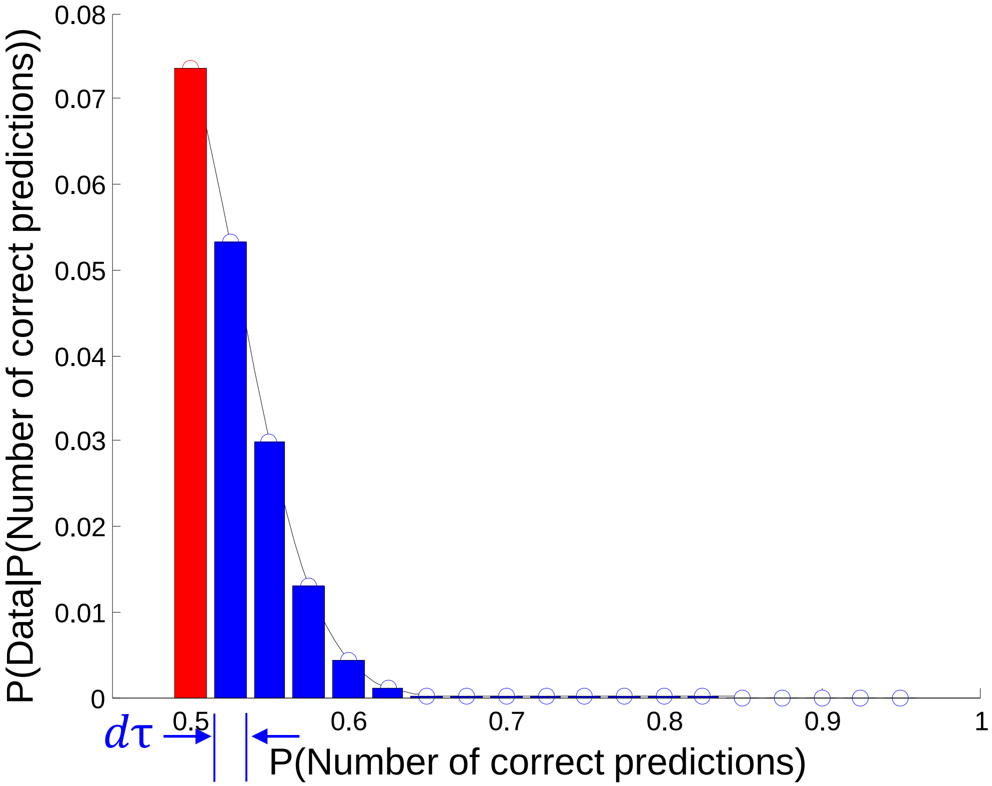

Example 1: Bayesian Inference

- In the end, the odds of the alternative hypothesis to the null hypothesis are 0.0930:1 (or roughly 1:10.75), indicating strong support for the null hypothesis.
- We conclude that it is most likely that the fortune teller is predicting stock price changes at chance.

Example 1: Summary

NHST

- Failed to reject the null hypothesis
- $P(\# \text{ correct} \geq 48 | P_{correct} = 0.5) = 0.6178$

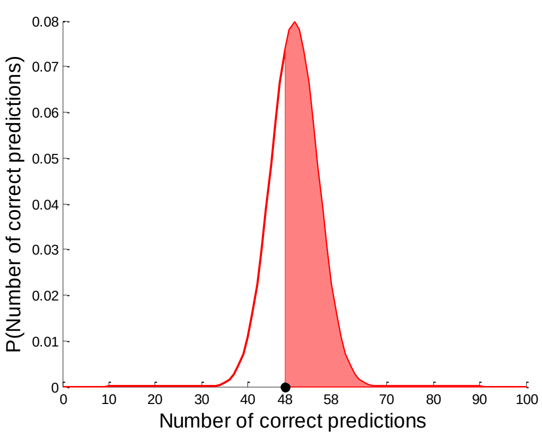

Bayes

- The observed data favours the null hypothesis over the alternative hypothesis.

$$\frac{P(P_{correct} > 0.5|48\ correct)}{P(P_{correct} = 0.5|48\ correct)} = 0.0930$$

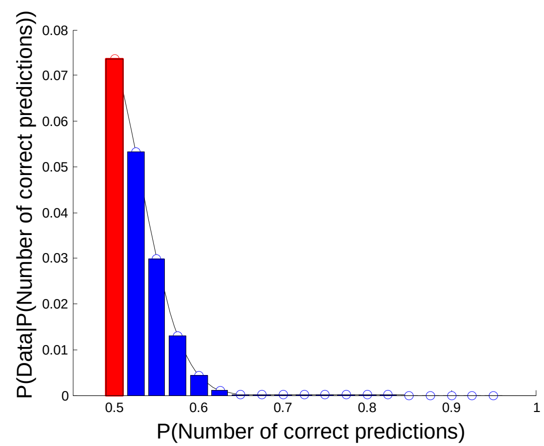

END Class 8b