Prior Knowledge is Power:

Bayesian Methods to Optimize Design in Clinical Development

ISPOR Workshop

May 6th 2024

Prior Knowledge is Power:

Bayesian Methods to Optimize Design in Clinical Development

Introduction to Bayesian Inference

- Craig Parzynski, Genesis Research Group

Regulatory Experience in Bayesian Clinical Trials

- Dr. Jennifer Clark, Food & Drug Administration

Overview of Bayesian Power Borrowing Methods in Clinical Trials

- Dr. Bret Zeldow, Genesis Research Group

Trial Design Considerations for Bayesian Borrowing

- Dr. Madhawa Saranadasa, Johnson & Johnson Innovative Medicine

Introduction to Bayesian Inference

Craig Parzynski

Disclaimer

The views expressed in this presentation are mine and not of Genesis Research Group. As an employee of Genesis Research Group, I am a paid consultant for Life Science companies.

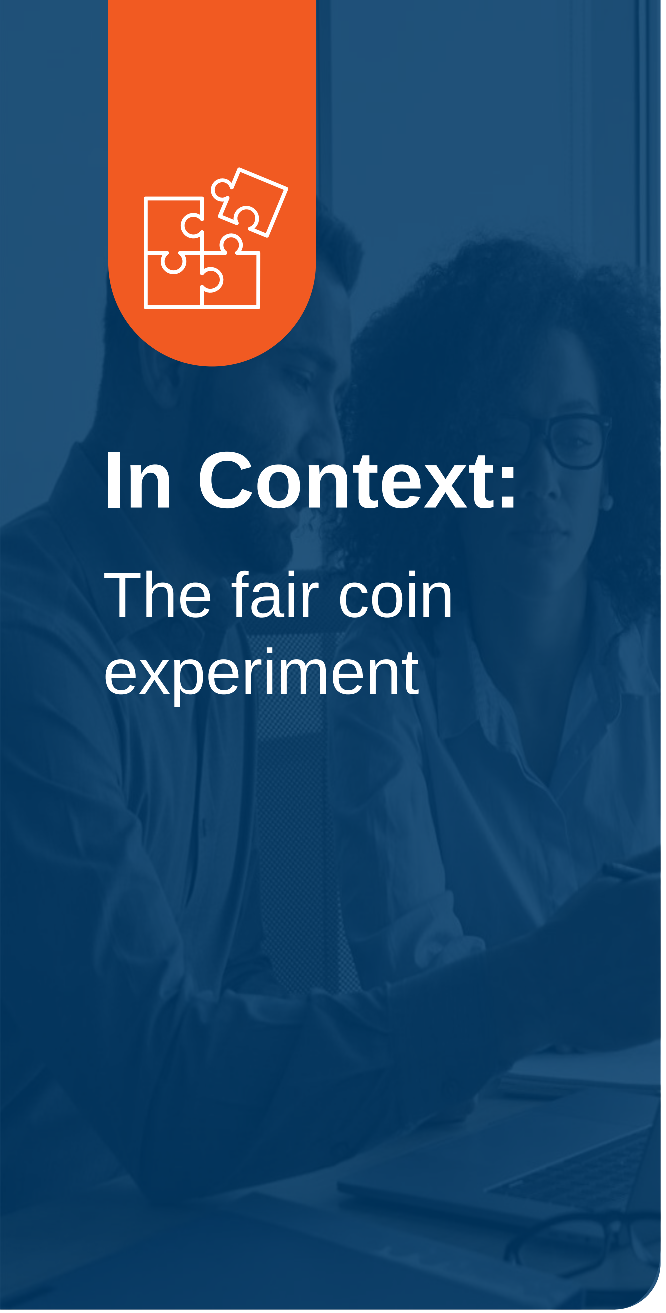

In Context:

The fair coin experiment

Frequentist	Bayesian
H0: p=.5 • Under the null our coin is fair H1: p >.5; • the coin favors heads	What's our prior belief about the coin? • We don't have any: Noninformative • It's more likely to be fair: Informed • Ask Pr(p > .5 | Data)
Run the experiment: Flip the coin 10x, Observe 7 heads	
Evaluate our data against H0 1. Statistical test: Z-test 2. p value: 0.103	Update our belief based on our new information 1. Estimate posterior distribution 2. Calculate Pr(p >.5 | Data): 0.86
Fail to reject H0, our coin could still be fair!	The probability our coin favors heads is .88

The nuts and bolts of Bayesian Inference

Bayes Theorem: The heart of Bayesian analysis

Conditional Probability

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

Start with $P(A)$. Use data to update to $P(A|B)$.

Example:

Let A be the event that someone you meet is from the Midwest.

- Using the US census: $P(A) = 0.20$

At dinner, you hear them order a "pop."
Let B denote the event that someone uses "pop."

Given data, we now update our prior:

$P(A|B) = 0.70$

GENESIS RESEARCH GROUP

Bayes Theorem: The heart of Bayesian analysis

Conditional Probability

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

Bayesian inference

$$P(\theta|x) = \frac{L(x|\theta)P(\theta)}{P(x)}$$

- θ: The quantity we want to estimate (mean, proportion, etc.)
- x: The data you observed
- $P(\theta)$: The prior distribution of θ
- $P(x)$: The marginal probability of x
- $L(x|\theta)$: The likelihood of observing x given θ
- $P(\theta|x)$: The posterior distribution of θ; a probability distribution

GENESIS RESEARCH GROUP

First some mathematical niceties: $P(\theta|x) = \frac{L(x|\theta)P(\theta)}{P(x)}$

- We mostly ignore $P(x)$
- It is a constant with respect to θ
- Because of this we often use the following formula for Bayesian inference:

Proportional to

$$P(\theta|x) \propto L(x|\theta)P(\theta)$$

GENESIS RESEARCH GROUP

The Likelihood: $P(\theta|x) \propto L(x|\theta)P(\theta)$

$L(x|\theta)$

The likelihood of observing x given θ

- The statistical model that we care about
  - A binomial likelihood for a binary outcome
  - A normal likelihood for a continuous outcome
  - Cox model
  - Logistic model
  - Poisson model

- Important Note: the likelihood is what is maximized to find the estimate of θ in frequentist methods but it does not consider the prior information!
  - The estimate is based solely on the model (likelihood) and the observed data

GENESIS RESEARCH GROUP

The Likelihood: $P(\theta|x) \propto L(x|\theta)P(\theta)$

$L(x|\theta)$

The likelihood of observing x given θ

- The statistical model that we care about
  - A binomial likelihood for a binary outcome
  - A normal likelihood for a continuous outcome
  - Cox model
  - Logistic model
  - Poisson model

- Important Note: the likelihood is what is maximized to find the estimate of θ in frequentist methods but it does not consider the prior information!
  - The estimate is based solely on the model (likelihood) and the observed data

The Prior: $P(\theta|x) \propto L(x|\theta)P(\theta)$

$P(\theta)$

Your prior belief about the distribution of θ

Selected by the researcher/statistician

- Can be informed:
  - Based on prior knowledge/data/experiences
  - E.g. historical trials, RWD
- Can be noninformative:
  - A flat distribution with large variance
  - E.g. Normal (0, 100000)

Caution: Informed priors with low variance can influence the posterior distribution heavily; especially in low sample situations

The Posterior Distribution: $P(\theta|x) \propto L(x|\theta)P(\theta)$

- The updated probability distribution of θ after collecting your data
- Used for inference in Bayesian statistics
- It is a weighted average of the prior information and your data

Choice of priors: Coin flip experiment

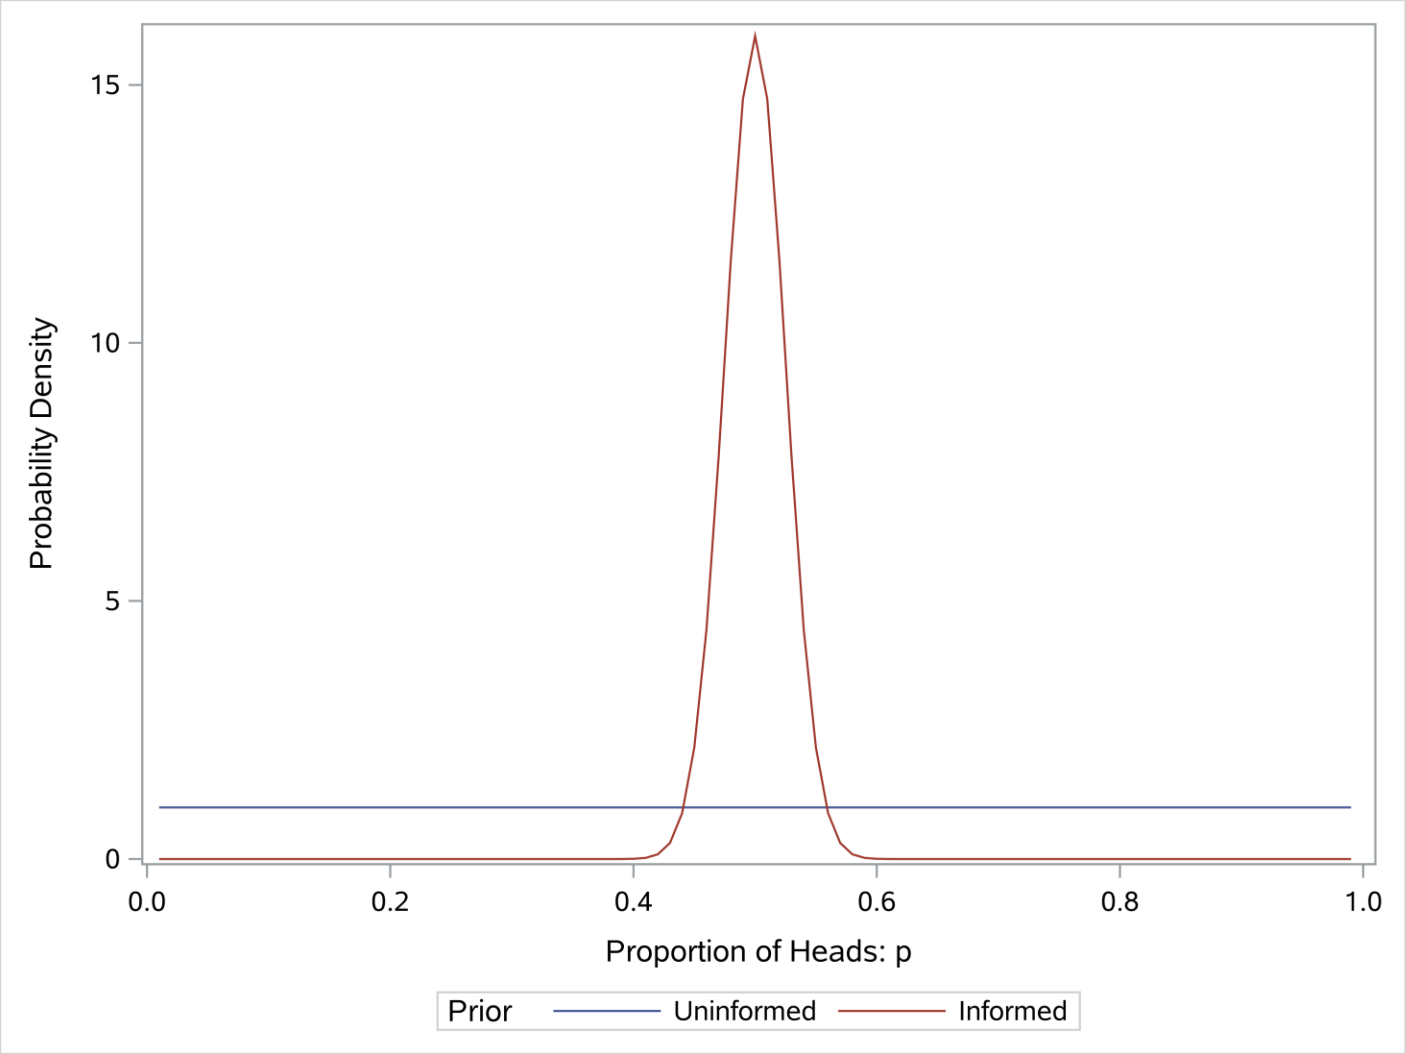

Choice of priors: Coin flip experiment

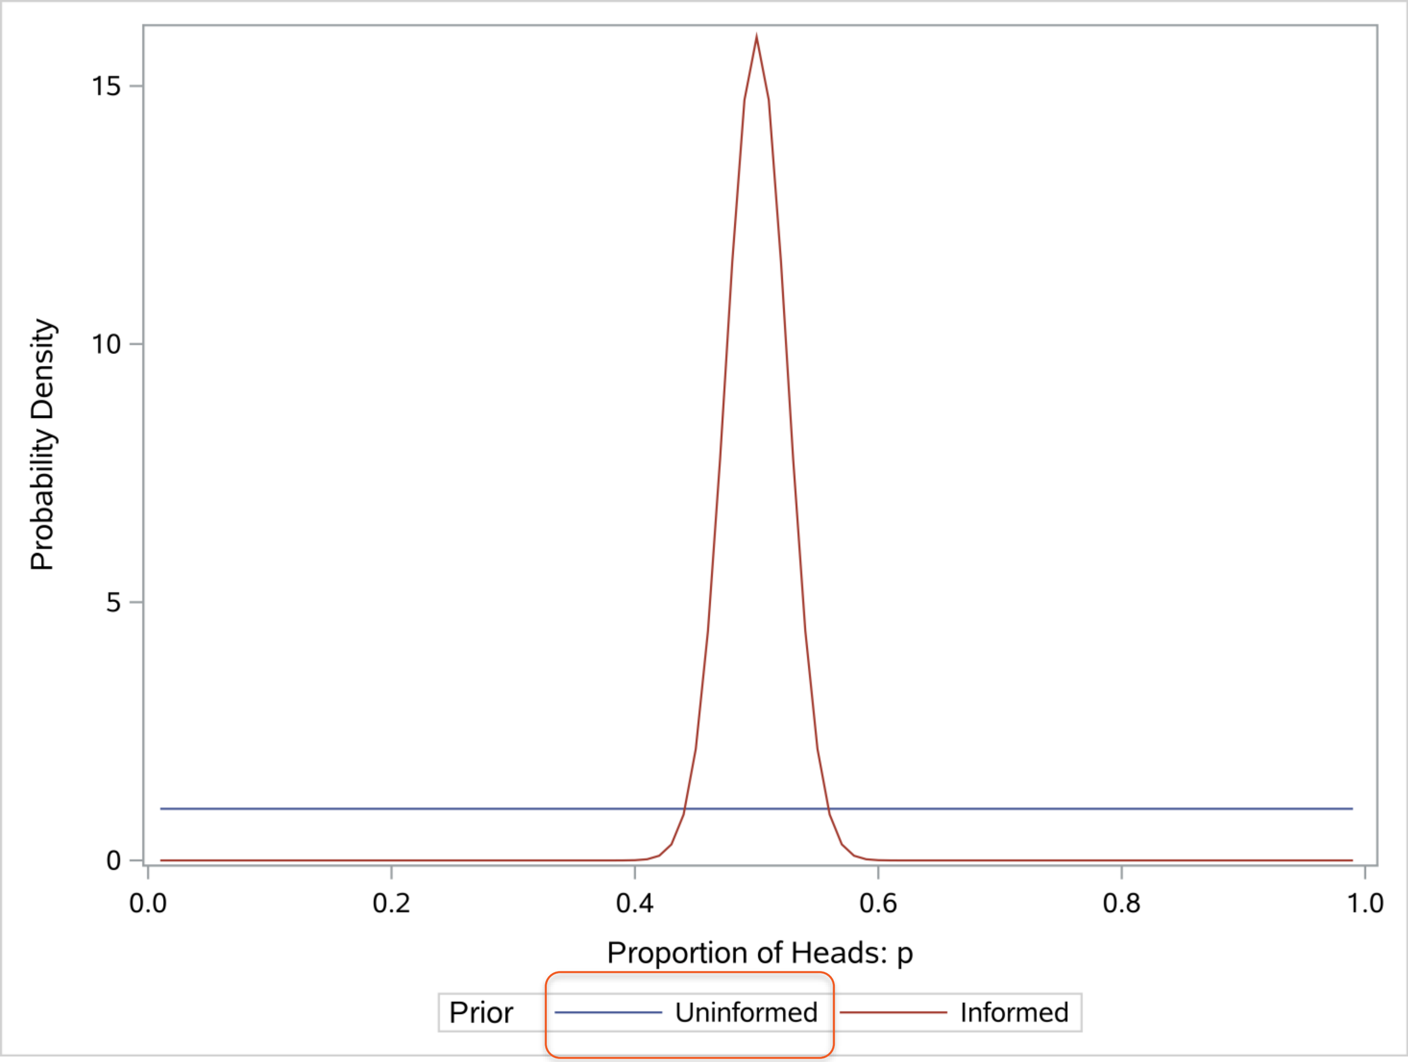

Posterior distribution using noninformative prior: Coin flip experiment

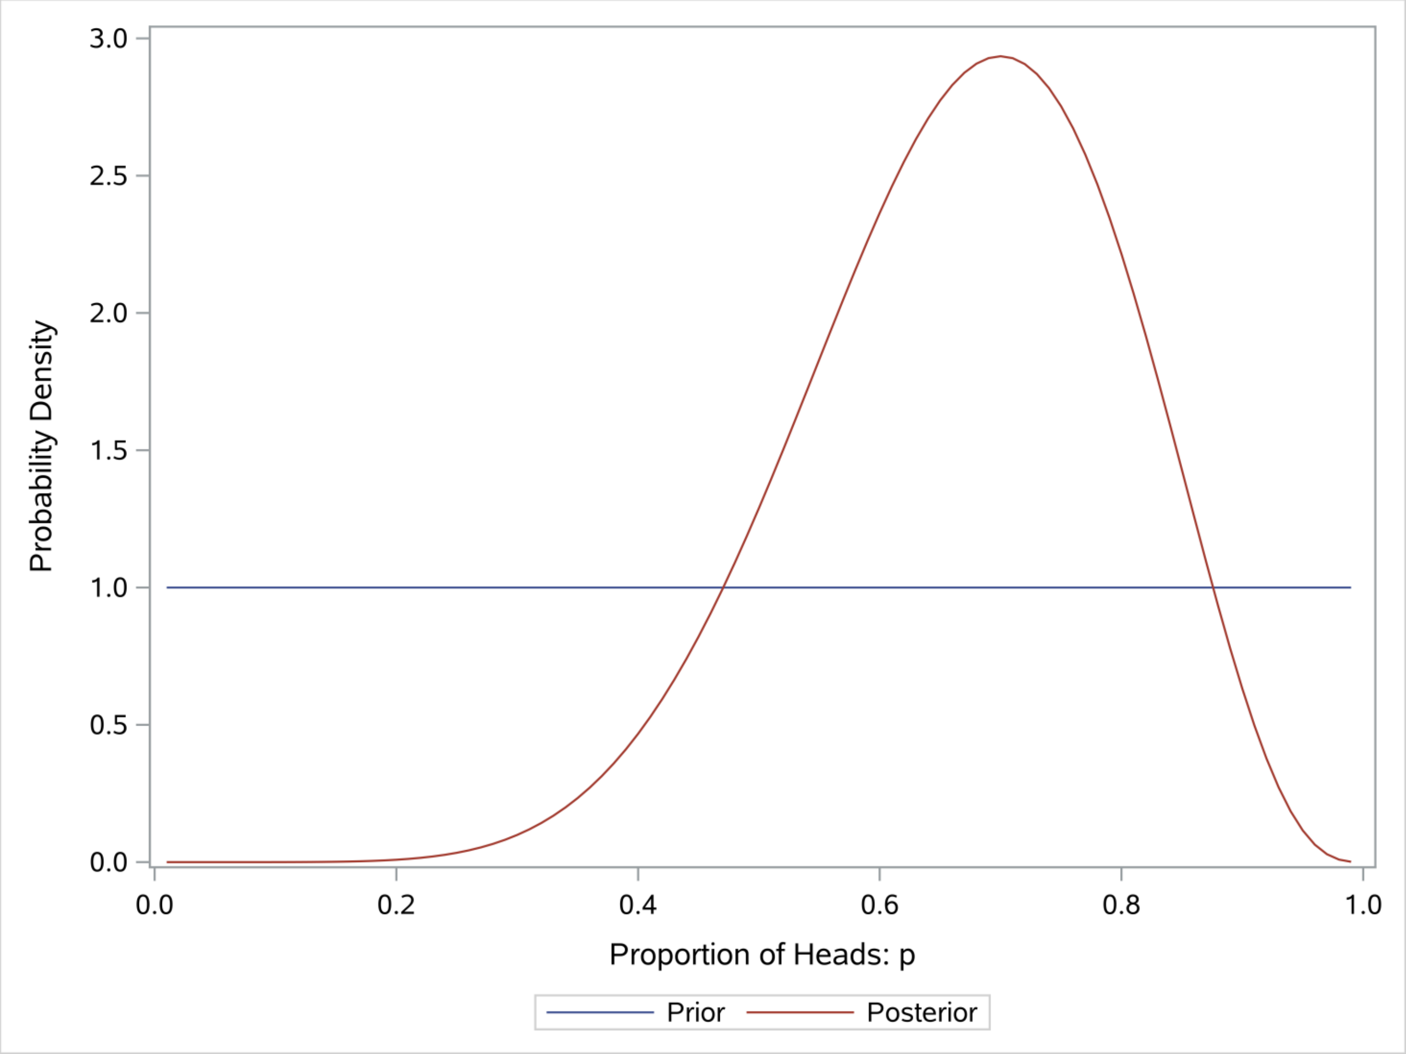

Posterior distribution using informative prior: Coin flip experiment

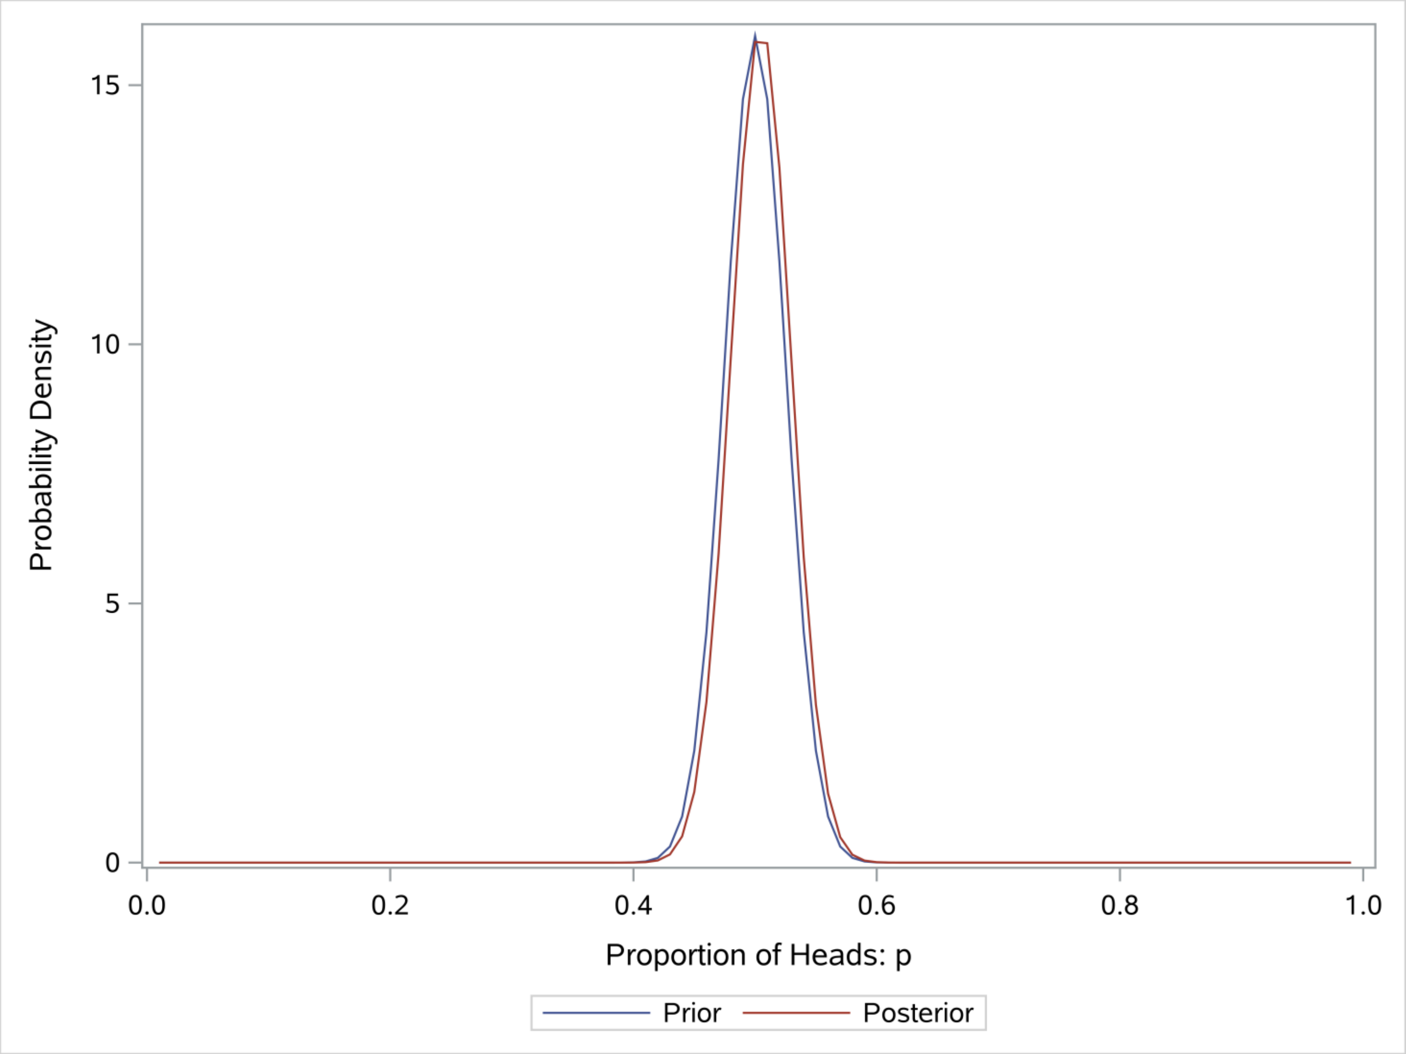

Posterior distributions: Coin flip experiment

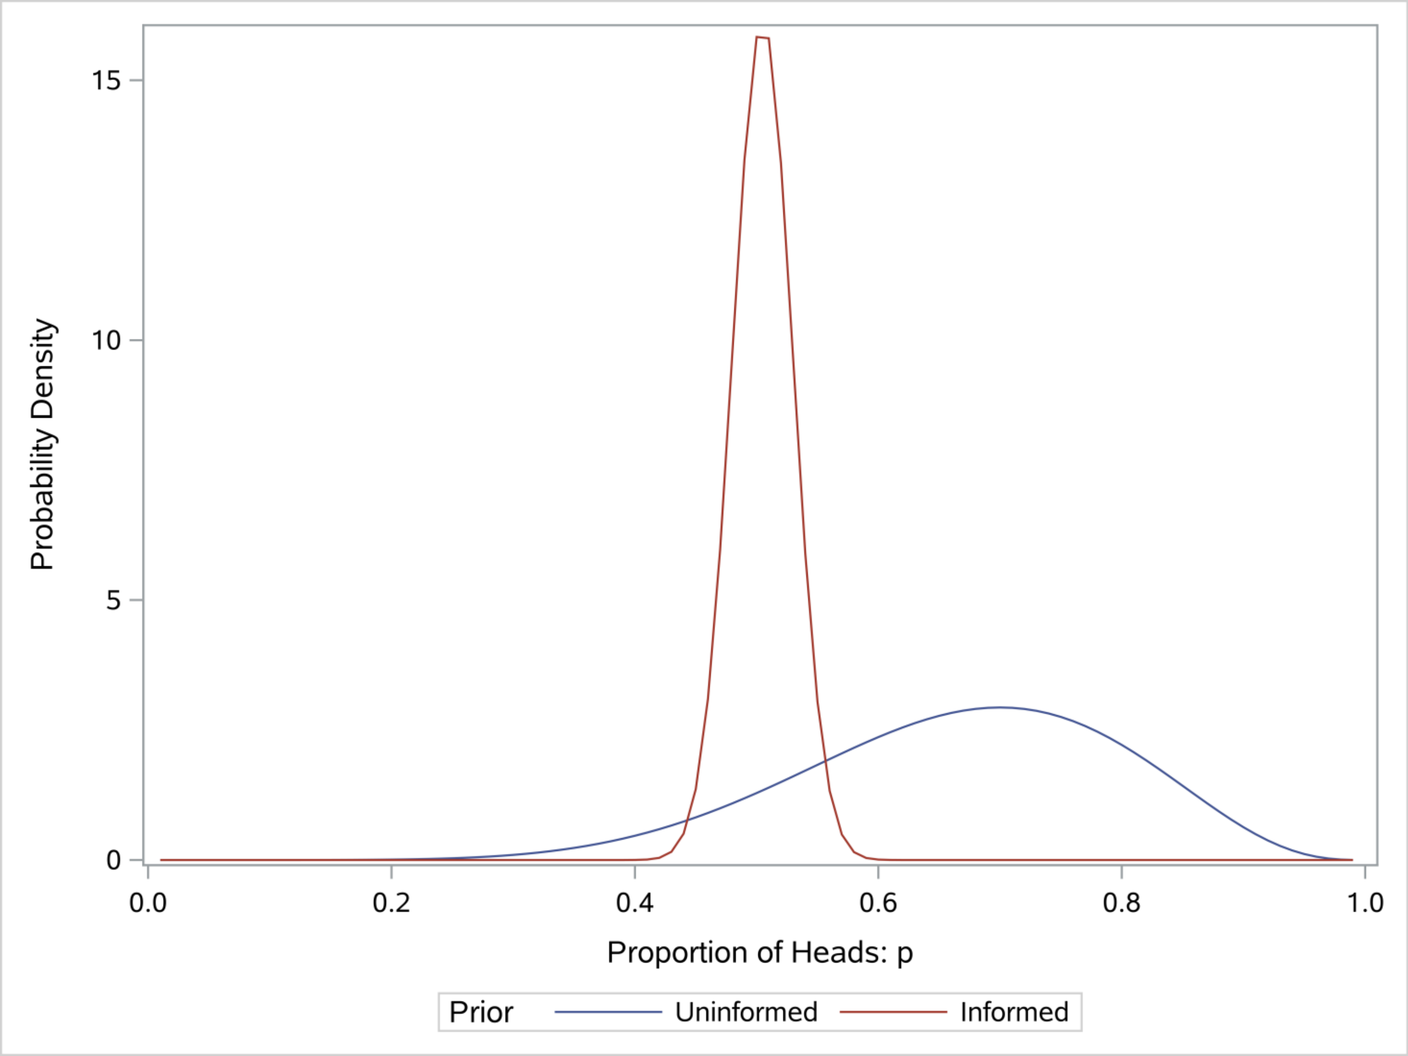

The Posterior Distribution: Inference

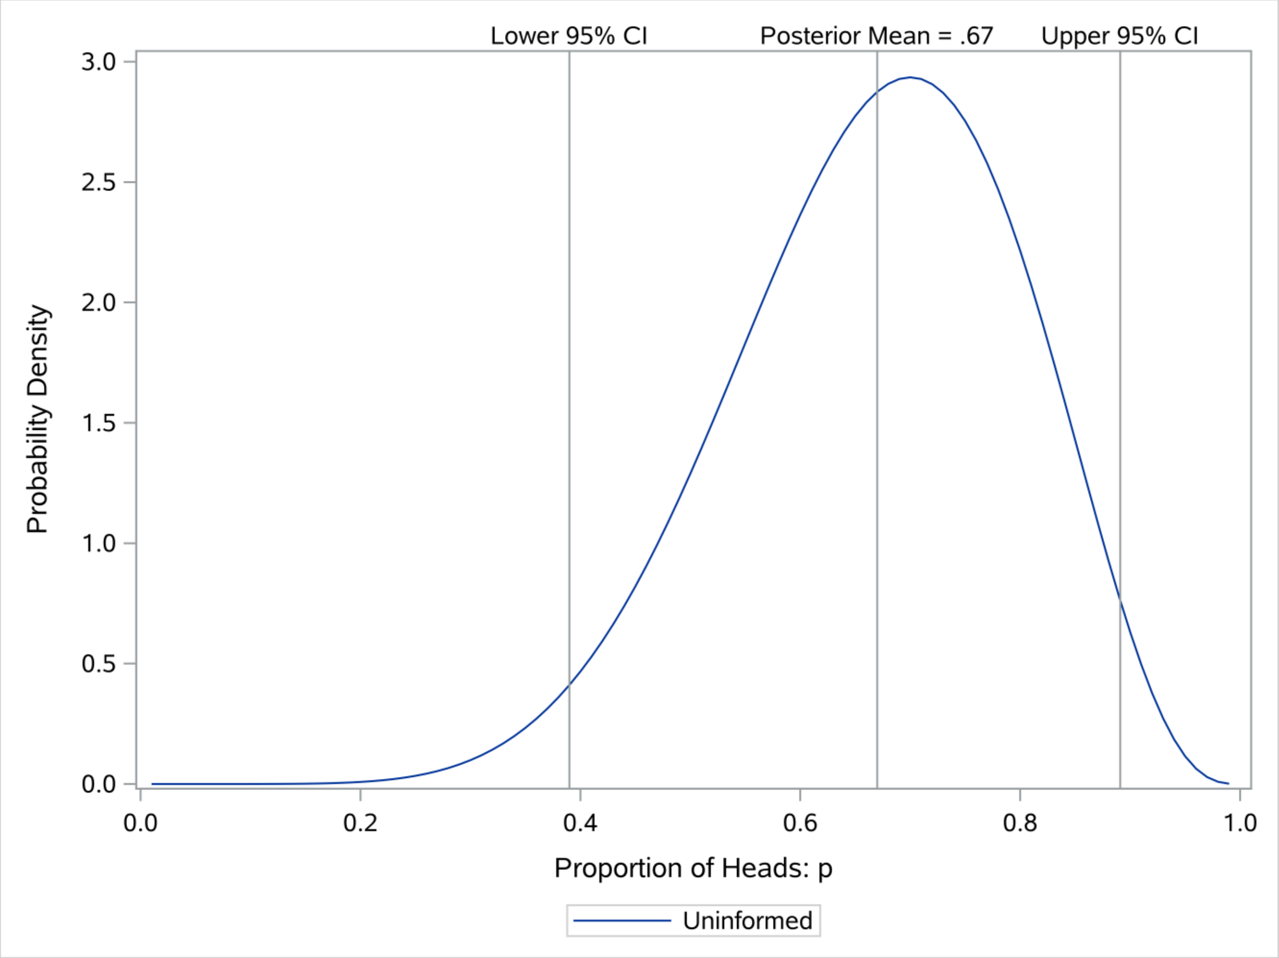

Posterior point estimate: mean, median

95% credible interval

- Interpreted as a 95% probability the mean is between two values
- Coin flip example:
  - Posterior mean (95% credible interval): 0.67 (0.39-0.88)

Probability statements (Bayesian p-value)

- Coin flip example:
  - Pr(p >.5 | Data): 0.88

GENESIS RESEARCH GROUP

The Posterior Distribution: Inference

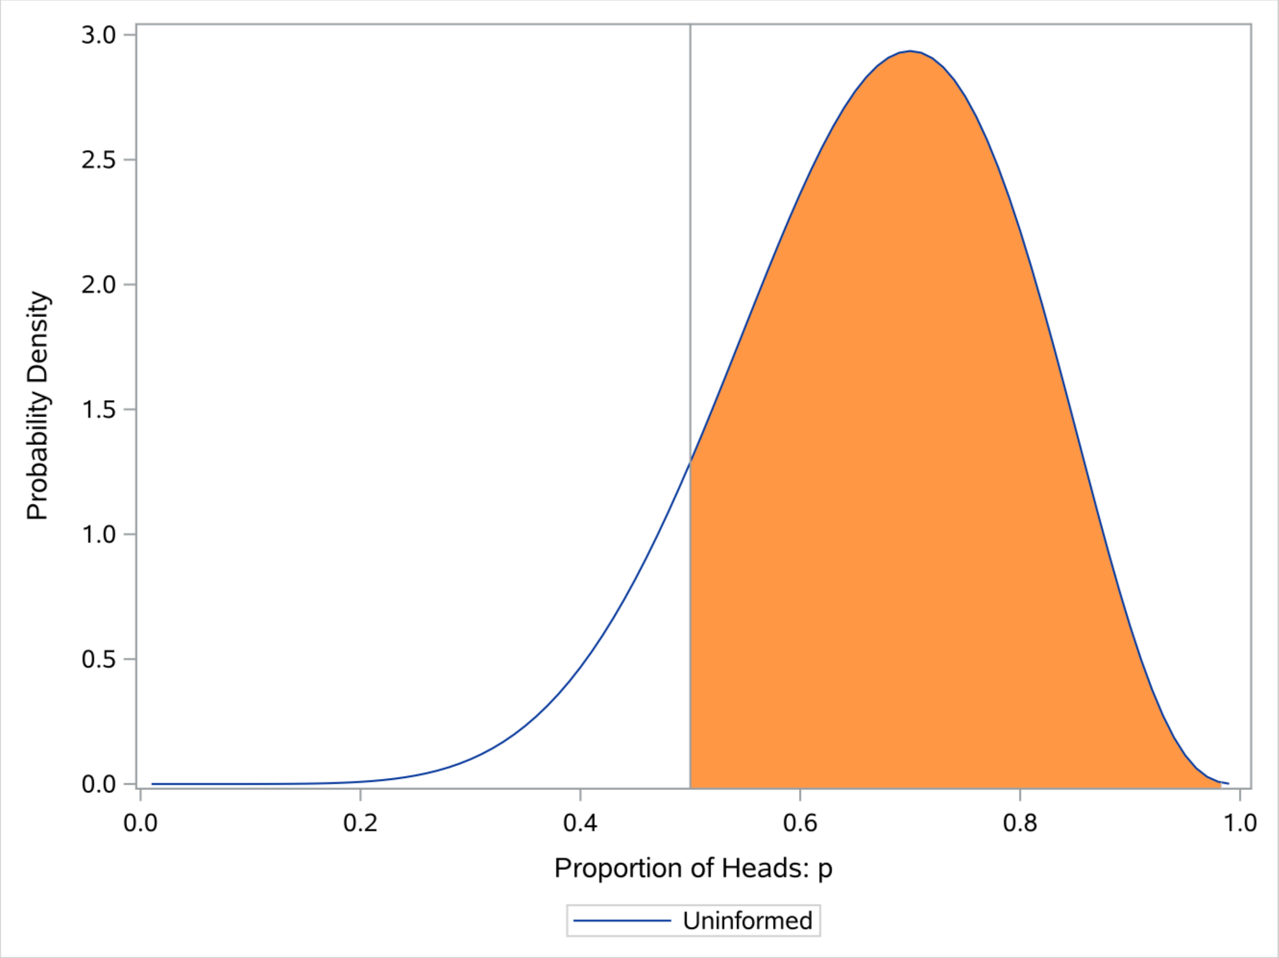

Posterior point estimate: mean, median

95% credible interval

- Interpreted as a 95% probability the mean is between two values
- Coin flip example:
  - Posterior mean (95% credible interval): 0.67 (0.39-0.88)

Probability statements (Bayesian p-value)

- Coin flip example:
  - Pr(p >.5 | Data): 0.88

To Summarize:

Bayesian inference uses **prior knowledge or information** to inform your parameter estimation

- Choice of your prior information can influence results

The **posterior distribution** is used to make inference and has the nice properties of being a probability distribution

- Posterior mean with 95% credible intervals
- Probability statements (Bayesian p-values)