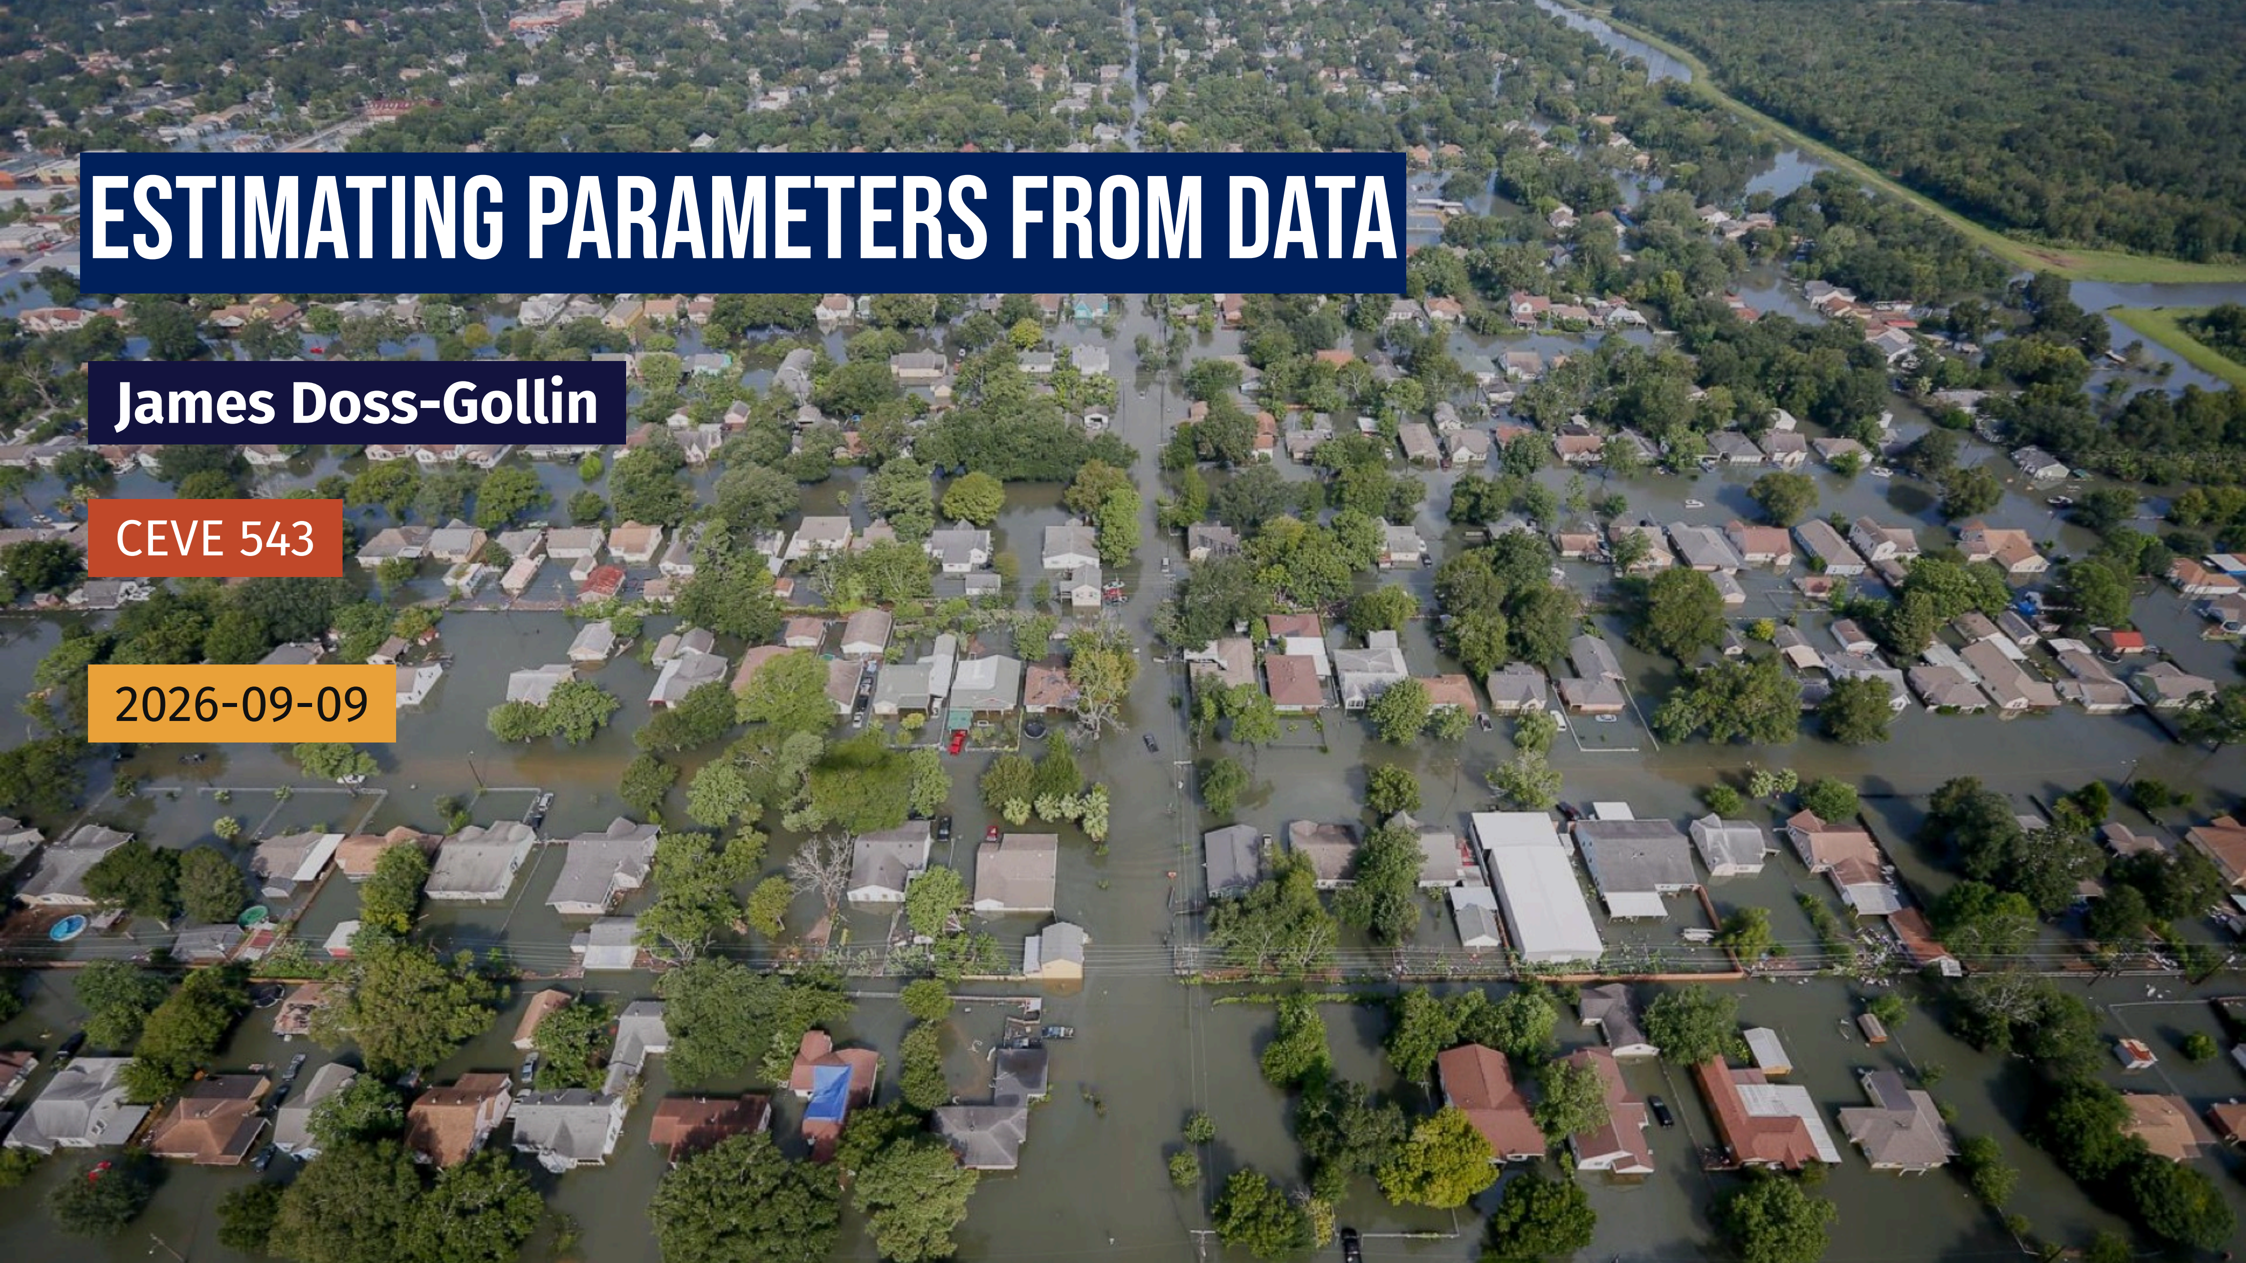

ESTIMATING PARAMETERS FROM DATA

James Doss-Gollin

CEVE 543

2026-09-09

Likelihood

Is a point estimate enough?

Bayesian estimation

HOW OFTEN DOES RAINFALL EXCEED THE CULVERT'S CAPACITY?

A stormwater culvert was sized to a design storm. You want to know how often rainfall exceeds that capacity: what is p, the annual probability of exceedance?

You have n years of record. Each year either exceeded the threshold ($Y_i = 1$) or did not ($Y_i = 0$). This is the coin flip from week 1, with unknown bias p. Here p plays the role of the θ from the board.

LIKELIHOOD OF p FROM 20 YEARS OF DATA

Suppose 3 of 20 years exceeded the threshold. The likelihood is the binomial probability of observing $k = 3$ in $n = 20$, viewed as a function of p:

$$L(p) = \binom{n}{k} p^k (1-p)^{n-k}$$

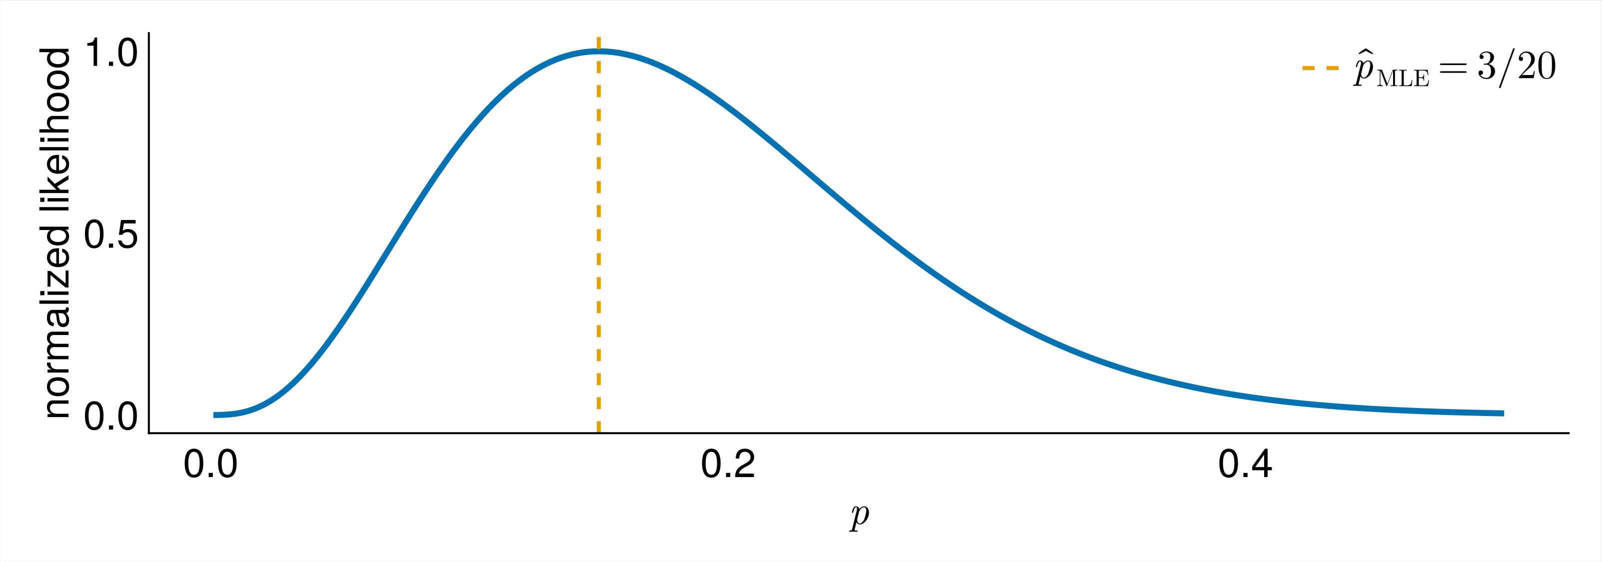

HOW THE LIKELIHOOD SHARPENS WITH MORE DATA

Same true $p = 0.15$, but now with 20, 100, and 500 years of record.

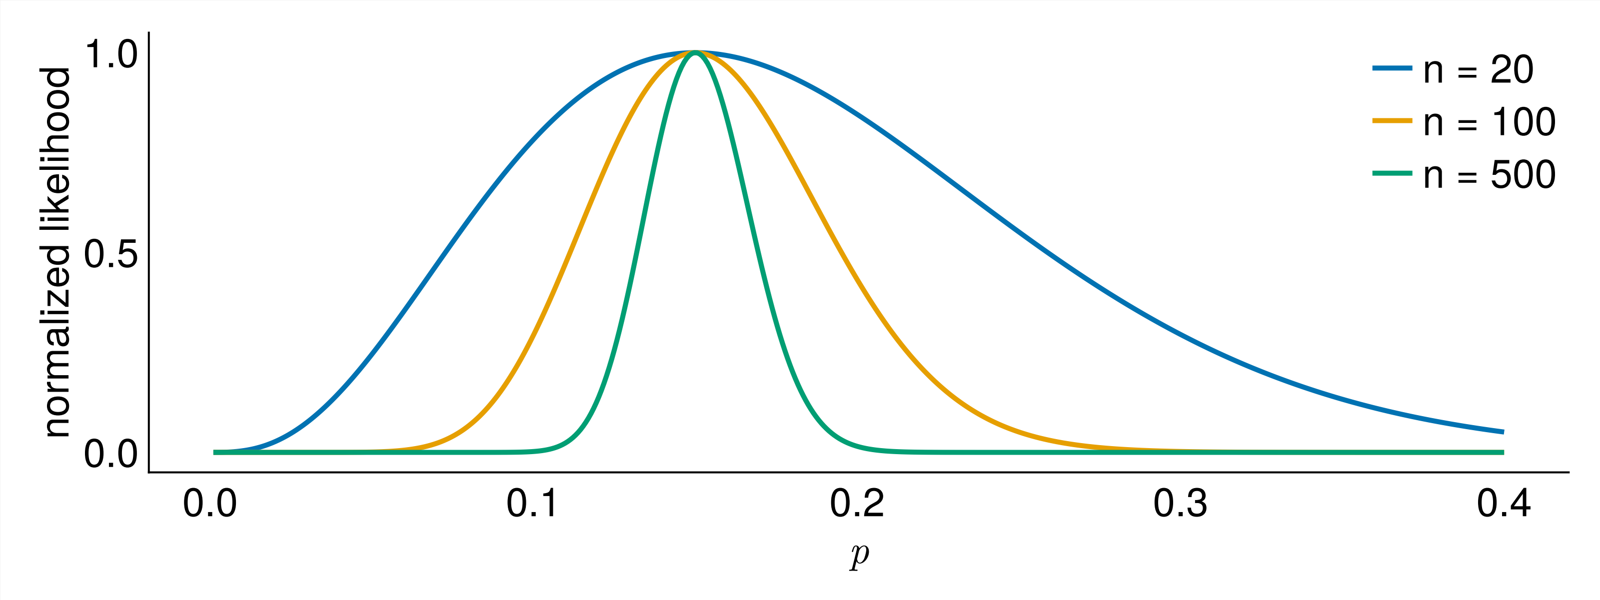

Even with 500 years of record the 95% CI for a roughly 7-year event is [0.12, 0.18]. Estimation is harder than it looks.

Likelihood

Is a point estimate enough?

Bayesian estimation

TWO ENGINEERS, SAME MLE

Both fit $\hat{p} = 0.15$. One has 20 years of record, the other has 200. Same point estimate, very different certainty.

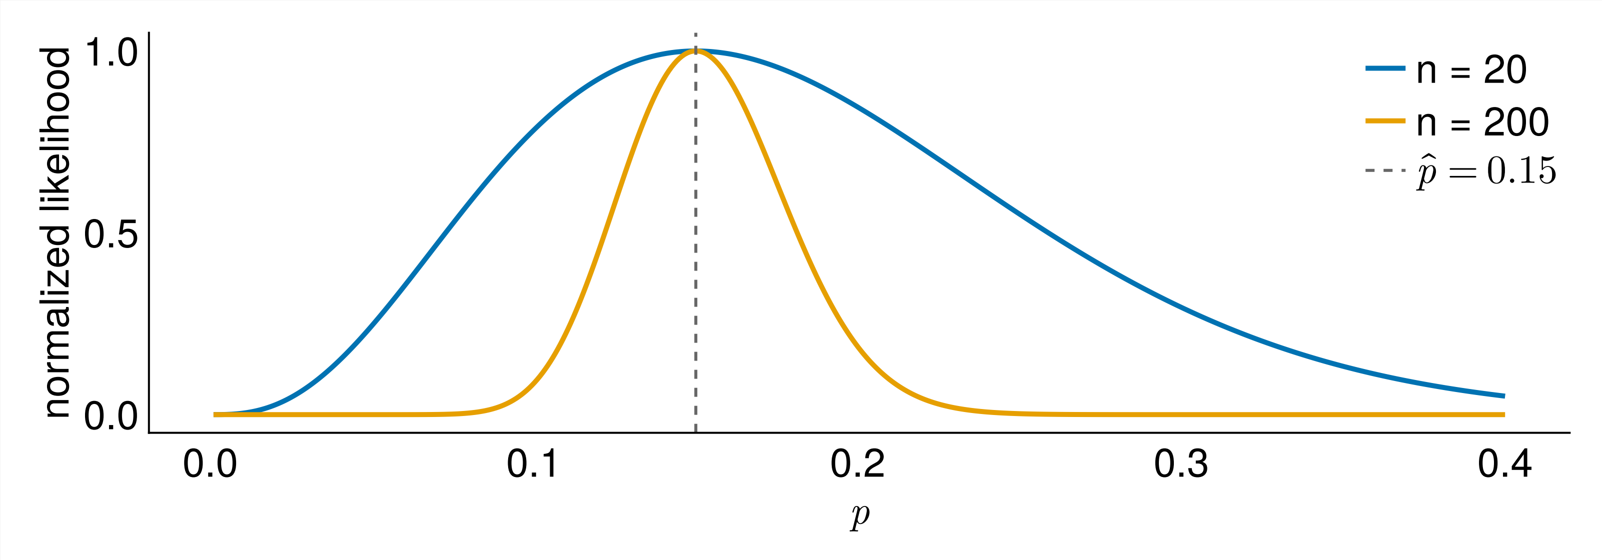

UNCERTAINTY FROM MLE: THE FISHER INFORMATION

The **Fisher information** $I(\theta)$ measures the curvature of the log-likelihood at the peak. Sharp peak → high information → low variance.

$$\mathrm{Var}(\hat{\theta}) \approx \frac{1}{I(\theta)}, \qquad \text{binomial: } \mathrm{Var}(\hat{p}) \approx \frac{p(1-p)}{n}$$

Tedious to compute for most models (not on the test). Bayes and the bootstrap give uncertainty more naturally.

Likelihood

Is a point estimate enough?

Bayesian estimation

BAYES' THEOREM IN ONE SLIDE

From the board:

$$\overbrace{p(\theta \mid y)}^{\text{posterior}} \propto \overbrace{p(y \mid \theta)}^{\text{likelihood}} \cdot \overbrace{p(\theta)}^{\text{prior}}$$

- **Prior** $p(\theta)$: what you believe before seeing data
- **Likelihood** $p(y \mid \theta)$: how consistent the data are with each θ
- **Posterior** $p(\theta \mid y)$: your updated belief, a full distribution over θ

THE BETA-BINOMIAL MODEL

The coin has unknown bias p. A Beta prior on p is **conjugate** to the Binomial likelihood, so the posterior has a closed form:

$$p \sim \mathrm{Beta}(\alpha, \beta), \quad Y \mid p \sim \mathrm{Binomial}(n, p) \quad \Rightarrow \quad p \mid Y = k \sim \mathrm{Beta}(\alpha + k, \beta + n - k)$$

$\alpha + \beta$ controls how strongly the prior concentrates. Most models are not conjugate; **Markov chain Monte Carlo** draws samples from the posterior when no closed form exists (week 5).

SMALL SAMPLE: THE PRIOR HELPS

True bias $p = 0.55$, only 10 flips, observed 7 heads. A Beta(10, 10) prior says coins are usually near 0.5.

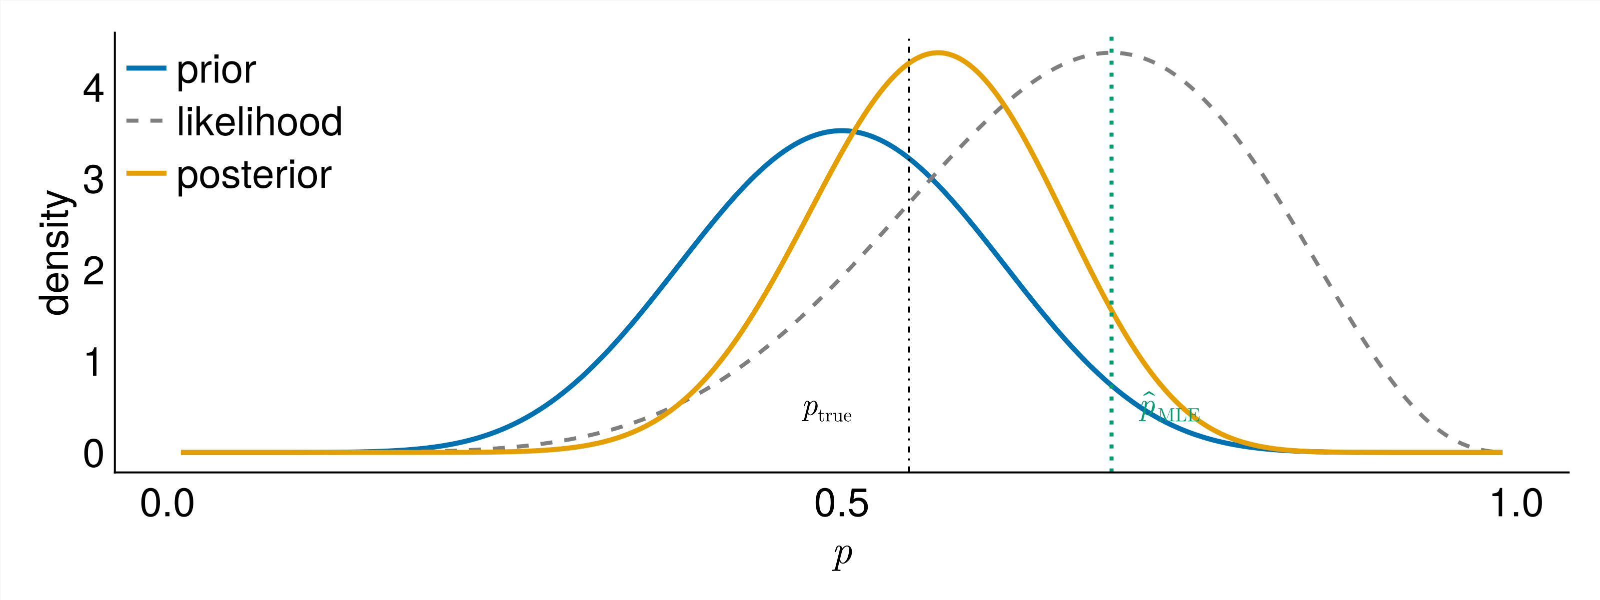

The MLE overshoots at 0.70. The posterior lands at 0.57, closer to the truth.

THE POSTERIOR IS A COMPROMISE BETWEEN PRIOR AND LIKELIHOOD

Same coin ($p = 0.55$), now with 10, 50, and 200 flips.

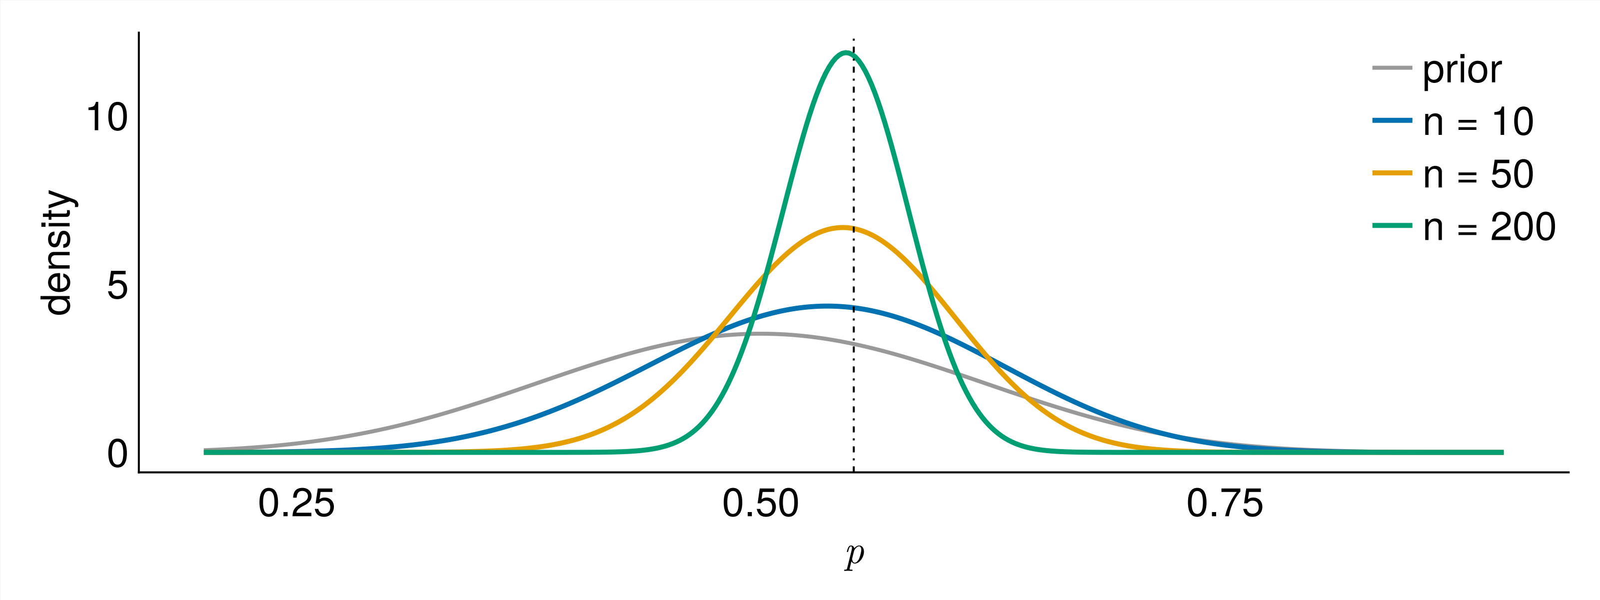

With enough data, two analysts with different priors converge.

THREE WAYS TO GET UNCERTAINTY

The posterior is a distribution over θ: credible interval, predictive distribution, threshold probability, all from the same object.

Method	Uncertainty	Needs
MLE	Asymptotic CI, per quantity	Model + Fisher information
Bayes	Posterior, all at once	Model + prior
Bootstrap	Resampled spread	Data only

WHAT IS NEXT

- **Friday**: the bootstrap and lab 3
- **Next week**: extreme value theory
- Module 1 test: **Friday October 16**