The Generalization Jigsaw

BAYESIANISM

- (methods)
- Parameter confidence
- Model choice: model weighting

FREQUENTISM

- (methods)
- Parameter confidence
- Model choice: hypothesis testing

EMPIRICISM

- Evaluating model fit
- Model choice: holdout evaluation

PUTTING THE JIGSAW TOGETHER

AN ILLUSTRATIVE THOUGHT EXPERIMENT

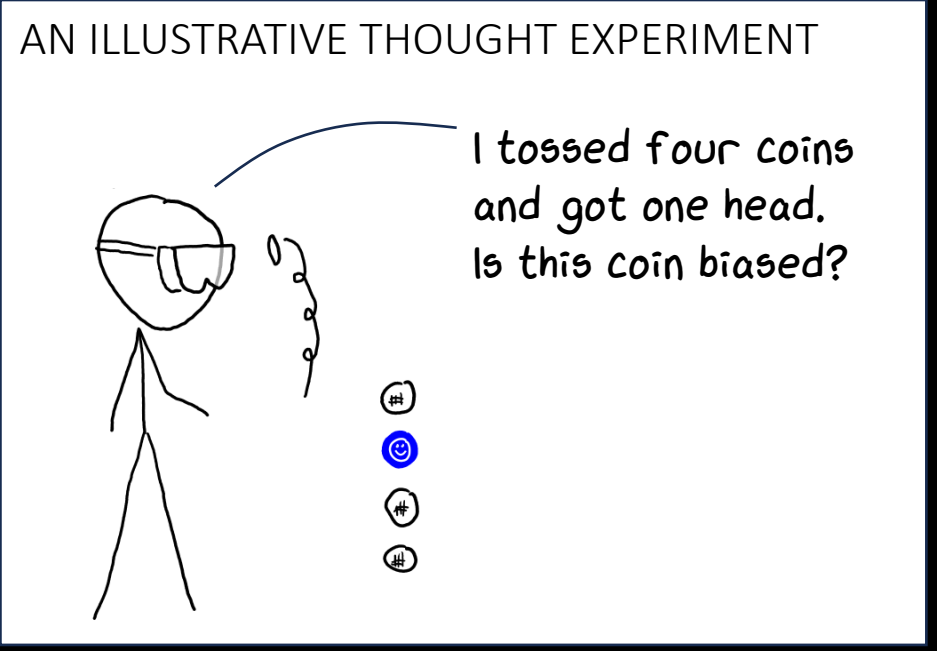

generalization
For this dataset, the probability of heads was p=25%. Can I conclude that the true value of p is 25%? No!

hypothesis testing
I could test the hypothesis that the coin is unbiased.

confidence intervals
Can I quantify how confident I am about p=25% ?

holdout set / model choice
I could test whether the p=25% model or the p=50% model fits better on a holdout dataset.

The Bayesianist approach to confidence intervals

By using random variables for unknown quantities,
we can reason about evidence and degree-of-belief.

probability of heads, unknown

Θ

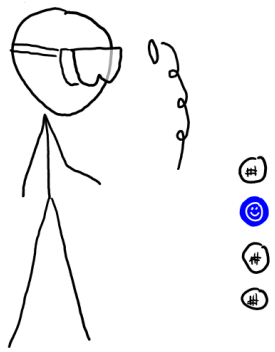

$X \sim \text{Bin}(4, \Theta)$

number of heads from 4 coin tosses

We don't know the value of Θ.
So let's treat it as a random variable.

We'll start by stating our prior belief.
Say, $\Theta \in \{0.25, 0.5, 0.75\}$, each equally likely.

We observed $X = 1$

We can use Bayes's rule to work out what this tells us about the unknown parameter's value …

$$P(\Theta = \tfrac{1}{2} \mid X = 1) = \frac{P(\Theta = \frac{1}{2}, X = 1)}{P(X = 1)} = \frac{P(\Theta = \frac{1}{2})\, P(X = 1 \mid \Theta = \frac{1}{2})}{P(X = 1)} = \frac{\frac{1}{3}\binom{4}{1}(\frac{1}{2})^1(\frac{1}{2})^3}{\sum_\Theta P(\Theta = \Theta)\, P(X = 1 \mid \Theta = \Theta)}$$

using the Law of Total Probability

prior belief

+

data

posterior belief

$\Theta \sim U[0,1]$

to express complete prior ignorance about θ

$X \sim \text{Bin}(n, \Theta)$

$$P(\Theta = 0.5 \mid X = 1) = \frac{P(\Theta = 0.5)\, P(X = 1 \mid \Theta = 0.5)}{P(X = 1)}$$

This is 0 when Θ is a cts r.v.

Even work when X is a cts r.v.

prior belief $\Pr_\Theta(\theta)$ + data x → posterior belief $\Pr_\Theta(\theta \mid X = x)$

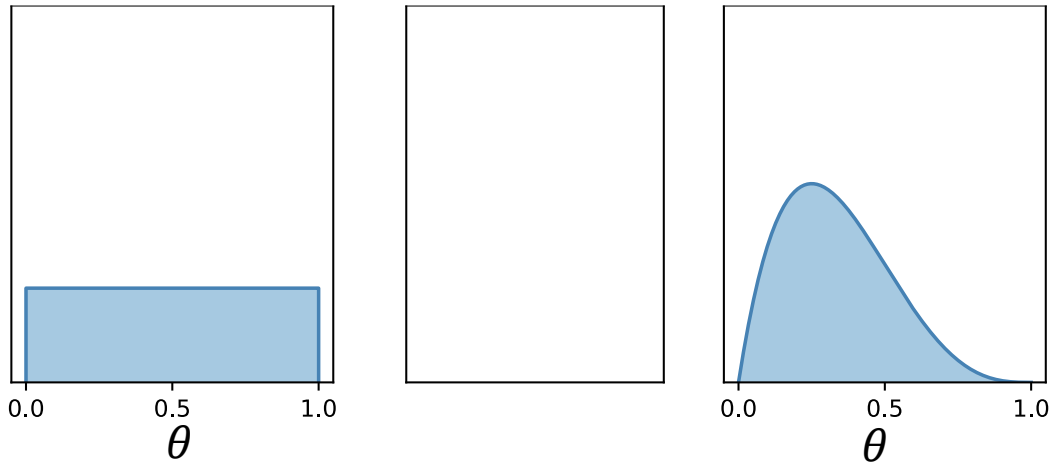

Reverend Thomas Bayes, 1701–1761

Bayes's rule for random variables

$$\mathbb{P}(X = x \mid Y = y) = \frac{\mathbb{P}(X = x)\,\mathbb{P}(Y = y \mid X = x)}{\mathbb{P}(Y = y)}$$

$$Pr_X(x \mid Y=y) = \frac{Pr_X(x)\; Pr_Y(y \mid X=x)}{Pr_Y(y)}$$

works when X or Y or both are continuous random variables

Bayesianism

Whenever there's an unknown parameter, you should express your uncertainty about it by treating it as a random variable.

Rationalist community

From Wikipedia, the free encyclopedia

> *This article is about the 21st-century movement. For the philosophical concept, see Rationalism. For other uses, see Rationalism (disambiguation).*

The **rationalist community** is a 21st-century movement that formed around a group of internet blogs, primarily LessWrong and *Astral Codex Ten* (formerly known as *Slate Star Codex*). The movement initially gained prominence in the San Francisco Bay Area. Its members seek to use rationality to avoid cognitive biases. Common interests include probability, effective altruism, transhumanism, and mitigating existential risk from artificial general intelligence.

The borders of the rationalist community are blurry and subject to debate among the community and adjacent groups.[1] Members who diverge from typical rationalist beliefs often self-describe as "rationalist-adjacent", "post-rationalist" (also known as "ingroup" and "TPOT", an acronym for "this part of Twitter"[2]) or "EA-adjacent".[3]

Description [edit]

Rationality [edit]

Rationalists define rationality to include epistemic rationality — coming to true beliefs about the world, and instrumental rationality — acting in a way to achieve one's objectives.[4]

The rationalists are concerned with applying science and probability to various topics,[5] with special attention to Bayesian inference.[6] According to Ellen Huet, the rationalist community "aim[s] to keep their thinking unbiased, even when the conclusions are scary".[7]