Summary

Given: an event E (usually n i.i.d. samples from a distribution with unknown parameter θ).

1. Find likelihood $\mathcal{L}(E;\theta)$

Usually $\prod \mathbb{P}\,(x_i;\theta)$ for discrete and $\prod f(x_i;\theta)$ for continuous

2. Maximize the likelihood. Usually:

A. Take the log (if it will make the math easier)

B. Set the derivative to 0 and solve

C. Use the second derivative test to confirm you have a maximizer

Expectation

$$\ln\left(\mathcal{L}\left(x_i;\theta_\mu,\theta_{\sigma^2}\right)\right) = \sum_{i=1}^{n} \ln\left(\frac{1}{\sqrt{\theta_{\sigma^2}2\pi}}\right) - \frac{1}{2}\cdot\frac{\left(x_i-\theta_\mu\right)^2}{\theta_{\sigma^2}}$$

$$\frac{\partial}{\partial\theta_\mu}\ln(\mathcal{L}) =$$

Setting equal to $\mathbf{0}$ and solving

Biased

One property we might want from an estimator is for it to be **unbiased**.

An estimator $\hat{\theta}$ is "unbiased" if
$$\mathbb{E}[\hat{\theta}] = \theta$$

The expectation is taken over the randomness in the samples we drew.

The formula is fixed, the data we draw to evaluate the formula becomes the source of the randomness.

So we're not consistently overestimating or underestimating.

If an estimator isn't unbiased then it's **biased**.

Are our MLEs biased?

Our estimate for the coin-flips (if we generalized a bit) would be $\frac{\text{num heads}}{\text{total flips}}$

pollev.com/robbie