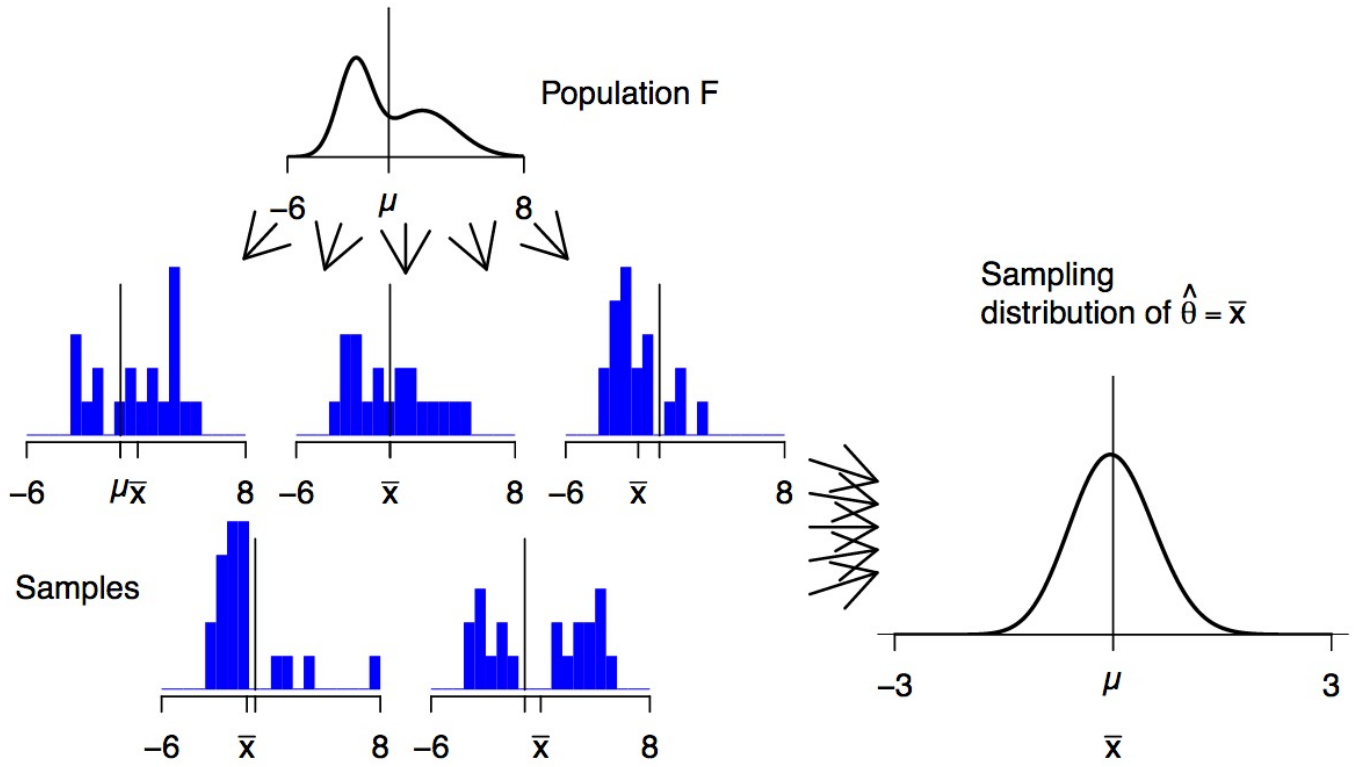


Figure 4: *Ideal world.* Sampling distributions are obtained by drawing repeated samples from the population, computing the statistic of interest for each, and collecting (an infinite number of) those statistics as the sampling distribution.

Estimating a population parameter by repeated samples from the population

Hesterberg 2014

https://arxiv.org/pdf/1411.5279.pdf

# Bootstrap distribution

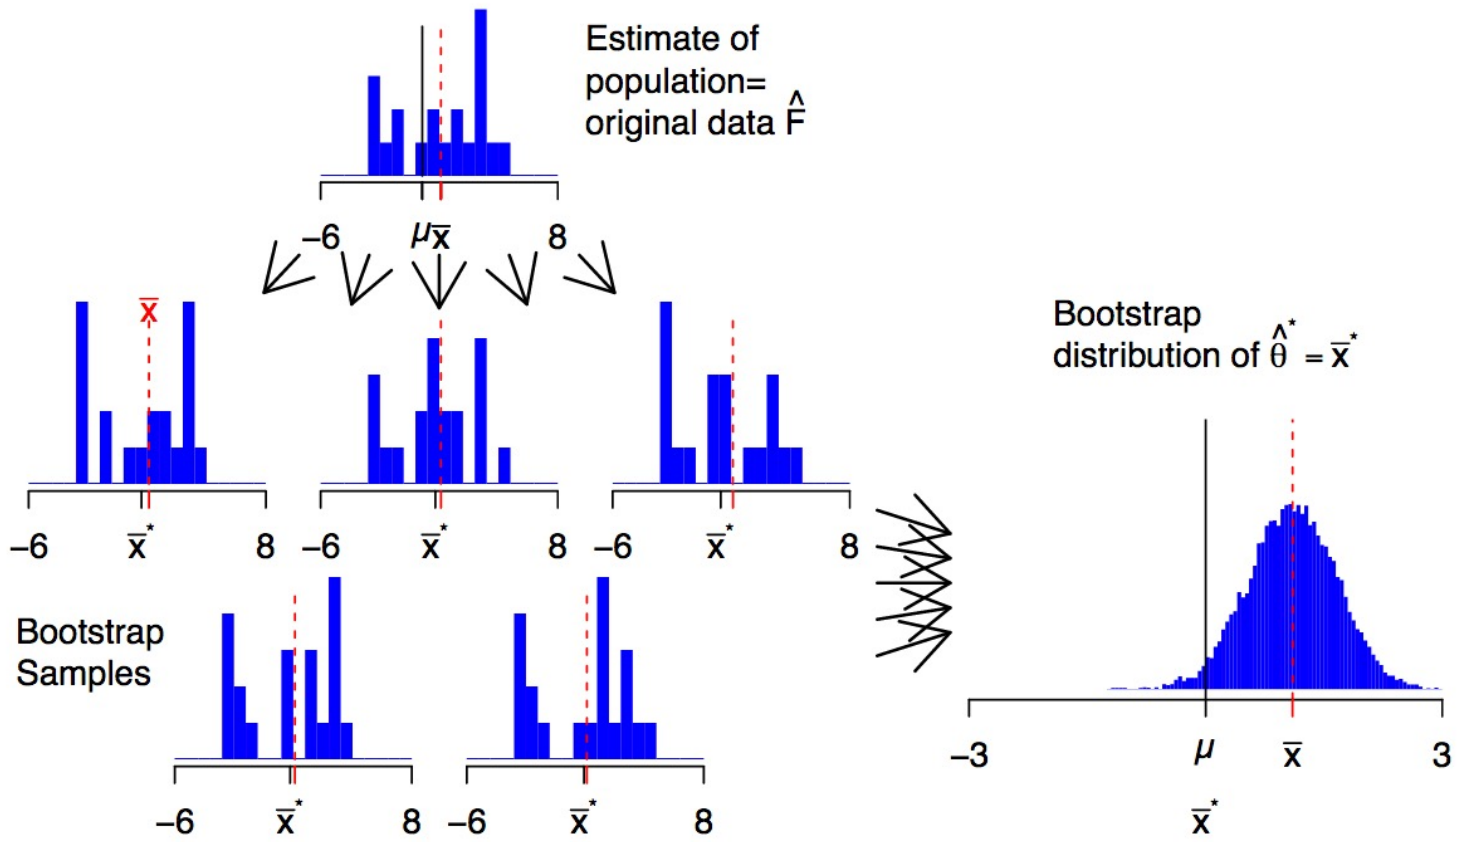


Figure 5: *Bootstrap world.* The bootstrap distribution is obtained by drawing repeated samples from an estimate of the population, computing the statistic of interest for each, and collecting those statistics. The distribution is centered at the observed statistic ($\bar{x}$), not the parameter ($\mu$).

Hesterberg 2014

https://arxiv.org/pdf/1411.5279.pdf

**Idea Behind the Bootstrap**

The idea behind the bootstrap is to estimate the population, then draw samples from that estimate, normally sampling the same way as in real life. The resulting *bootstrap distribution* is an estimate of the sampling distribution.

We use this for inferences, not to obtain better estimates. It is centered at the statistic (e.g. $\bar{x}$) not the parameter ($\mu$).

**Inference, Not Better Estimates** *The bootstrap distribution is centered at the observed statistic, not the population parameter*, e.g. at $\bar{x}$, not $\mu$.

This has two profound implications. First, it means that we do not use the bootstrap to get better estimates[4]. For example, we cannot use the bootstrap to improve on $\bar{x}$; no matter how many bootstrap samples we take, they are always centered at $\bar{x}$, not $\mu$. We'd just be adding random noise to $\bar{x}$. Instead we use the bootstrap to tell how accurate the original estimate is.

Hesterberg 2014

https://arxiv.org/pdf/1411.5279.pdf

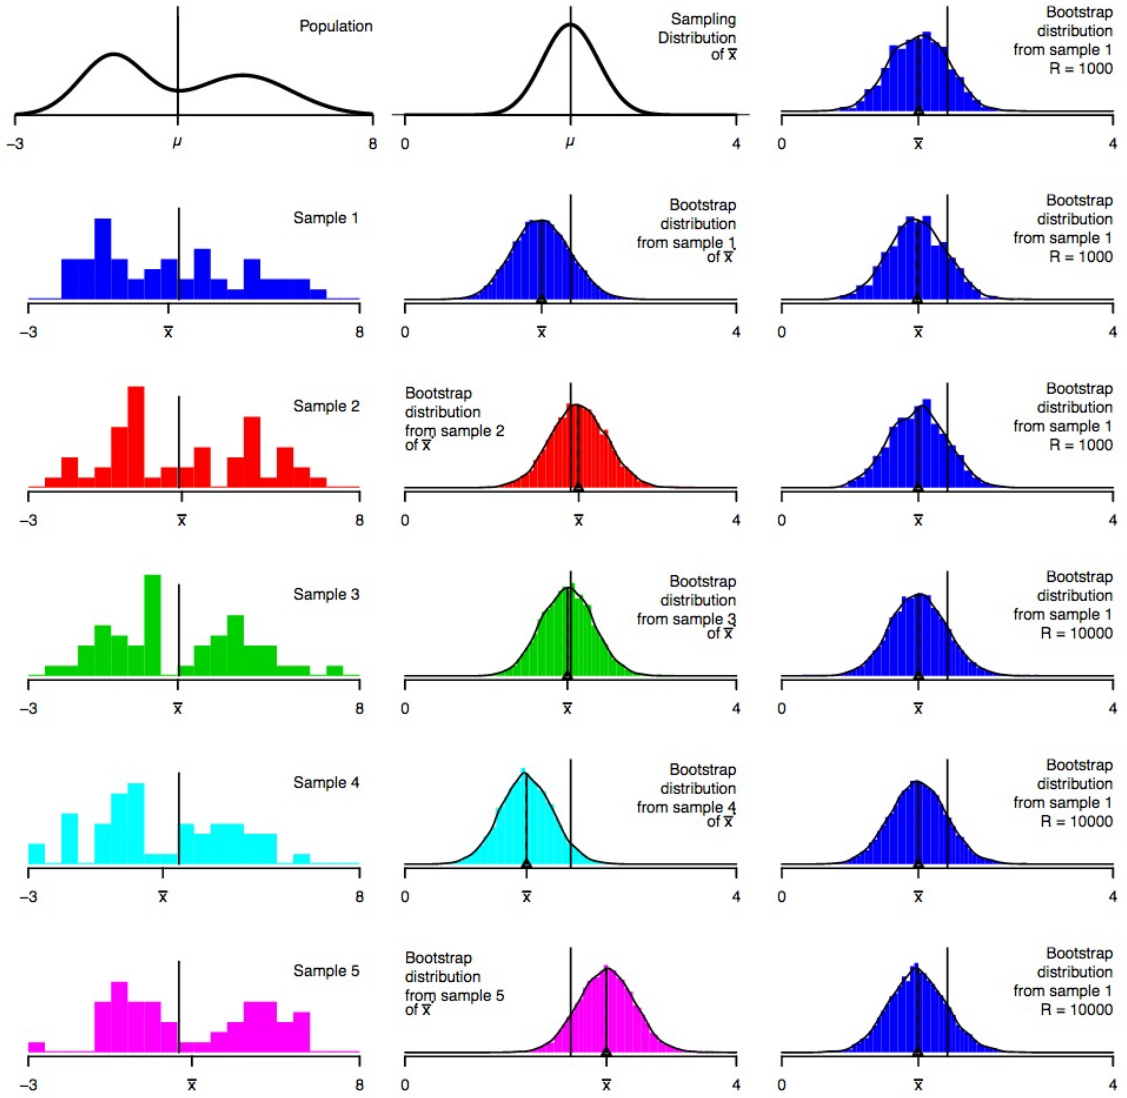


Figure 6: *Bootstrap distribution for the mean,* $n = 50$. The left column shows the population and five samples. The middle column shows the sampling distribution for $\bar{X}$, and bootstrap distributions of $\bar{X}^*$ from each sample, with $r = 10^4$. The right column shows more bootstrap distributions from the first sample, three with $r = 1000$ and three with $r = 10^4$.

Hesterberg 2014

https://arxiv.org/pdf/1411.5279.pdf

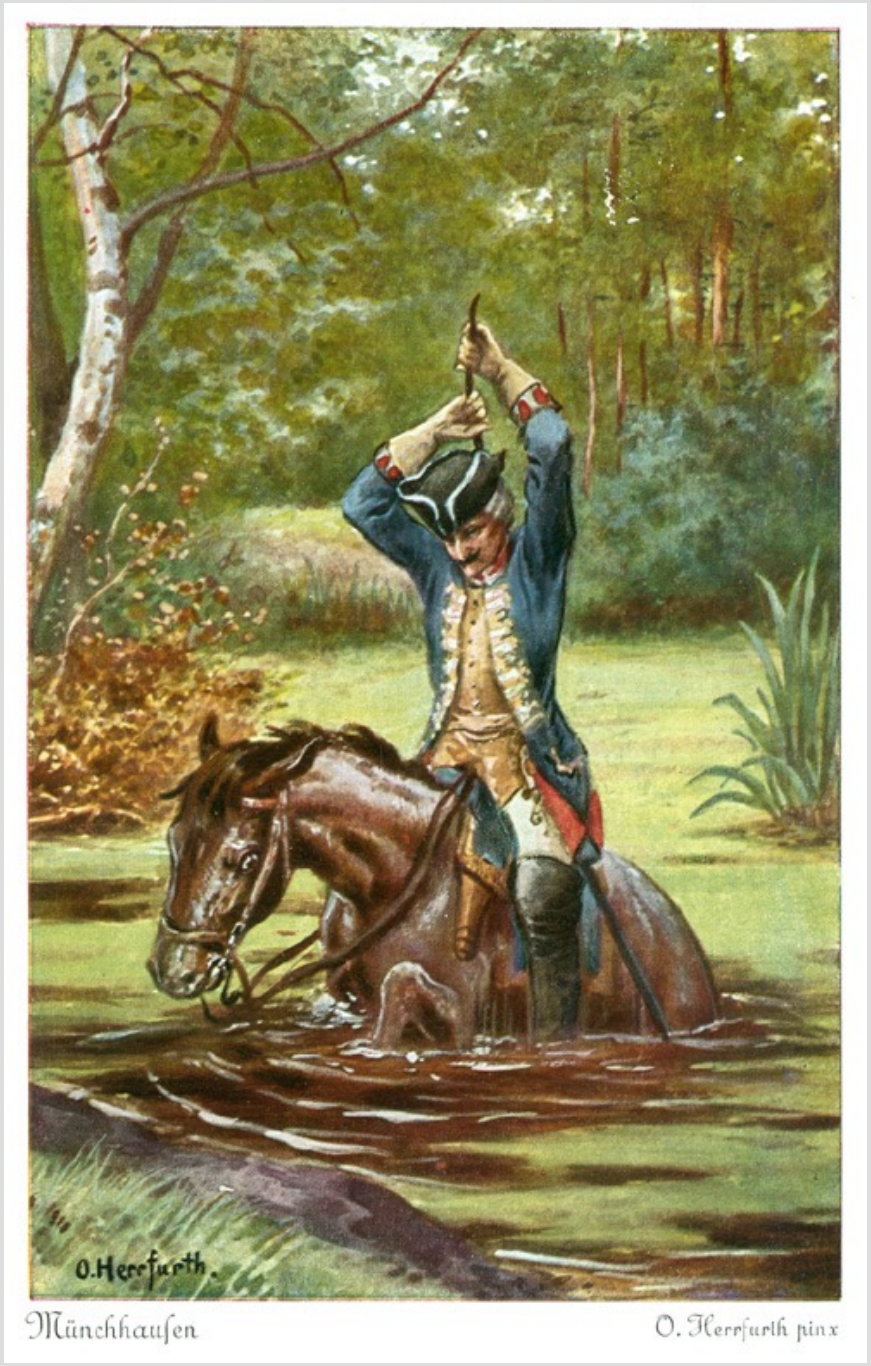


Baron Münchhausen has (in English version) been said to have pulled himself up from a swamp by his bootstrap

(To spoil our connection, in German, he did this by his hair.)

In any case, ”bootstrap” was named as such in statistics (likely) to represent the idea of using one’s existing data to ”pull out seemingly more data”, reminding of the miraculous trick of the Baron.

In practice, of course, nothing miraculous is happening here and we are not getting new data from the existing data. Instead, what we can pull out are some inferences about the accuracy of the estimator based on the existing data.

https://en.wikipedia.org/wiki/M%C3%BCnchhausen_trilemma#/media/File:Muenchhausen_Herrfurth_7_500x789.jpg

# 5.2 THE BOOTSTRAP

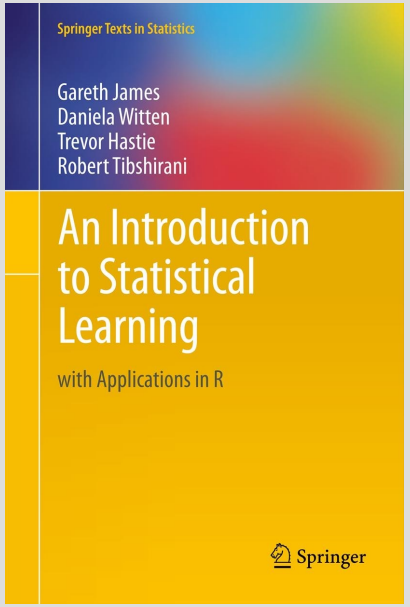


https://www.statlearning.com/