Agenda

Multi-parameter posterior

1. Administrative
2. Modeling income with two parameters
3. MAP estimation
4. ***Hands on:*** *Estimation in R*

Slides are licensed under CC BY-NC-SA 4.0
Uploading to generative AI tools or use for training machine learning models is prohibited

Lecture recordings

- Recordings of lectures are now available in the “recordings” folder on the fileshare

Modeling income

Annual family income in Canada, 2021

(weighted subsample of employed; n=3195)

Data source:
Canadian Income Survey: Public Use Microdata

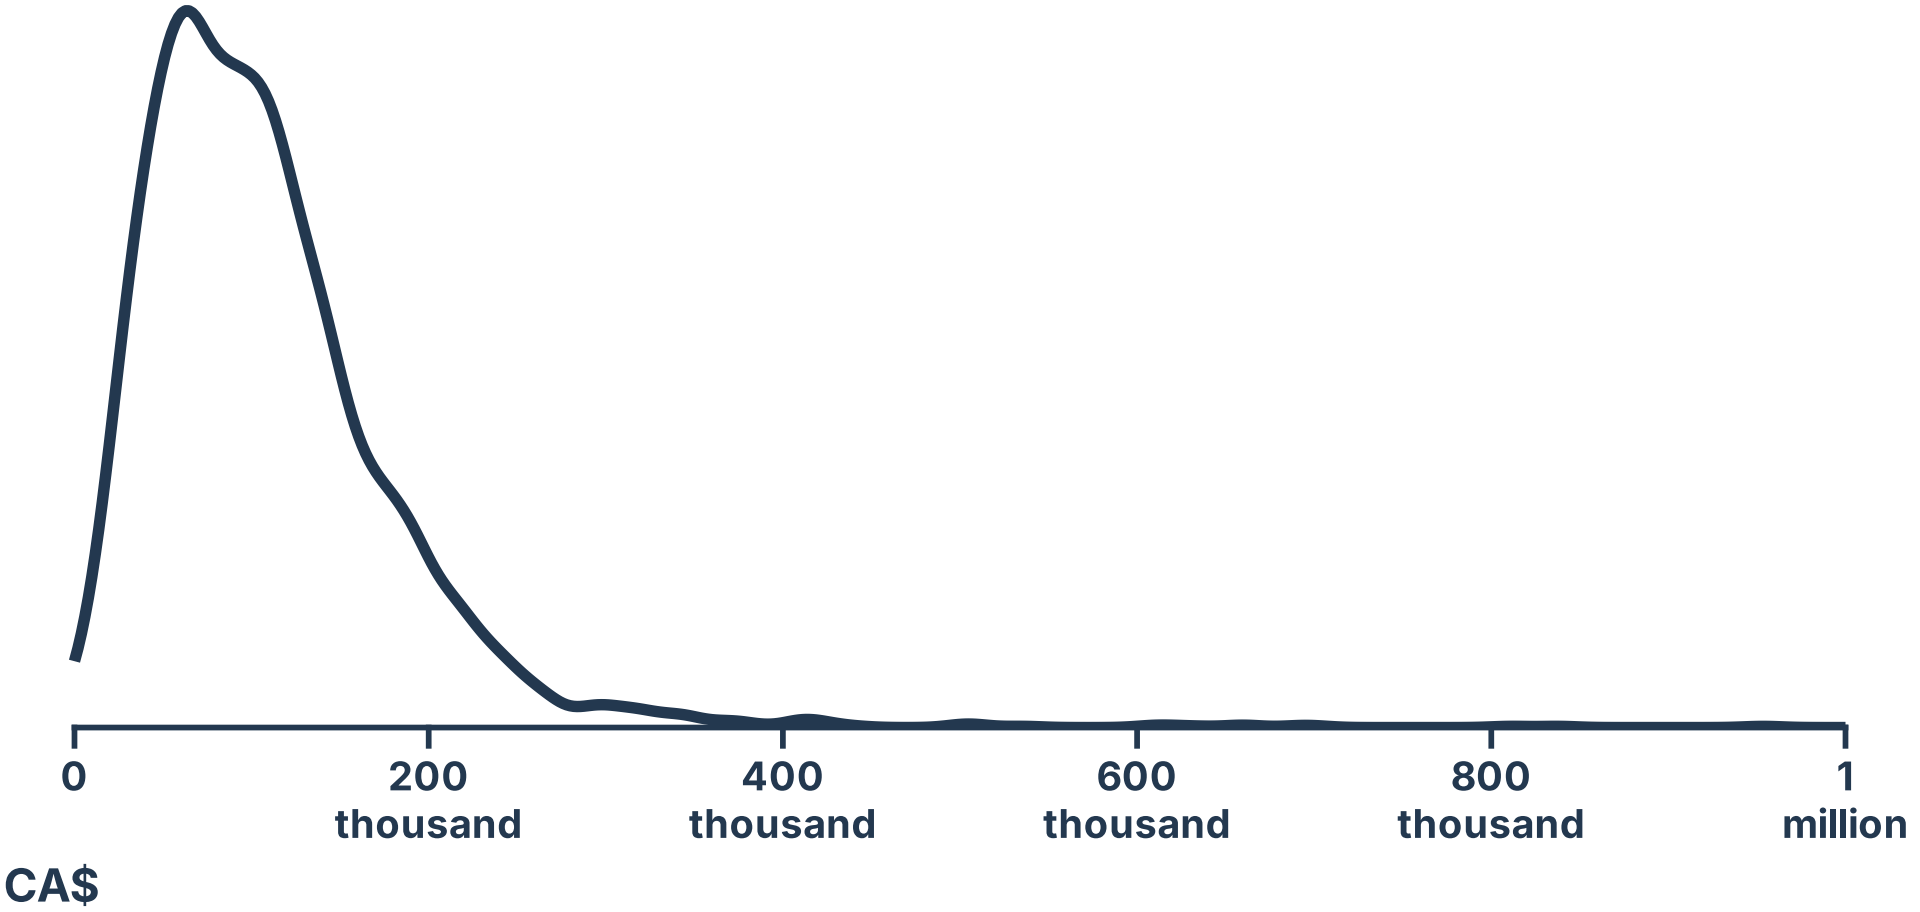

Annual family income in Canada, 2021

(weighted subsample of employed; n=3195)

Data source:
Canadian Income Survey: Public Use Microdata

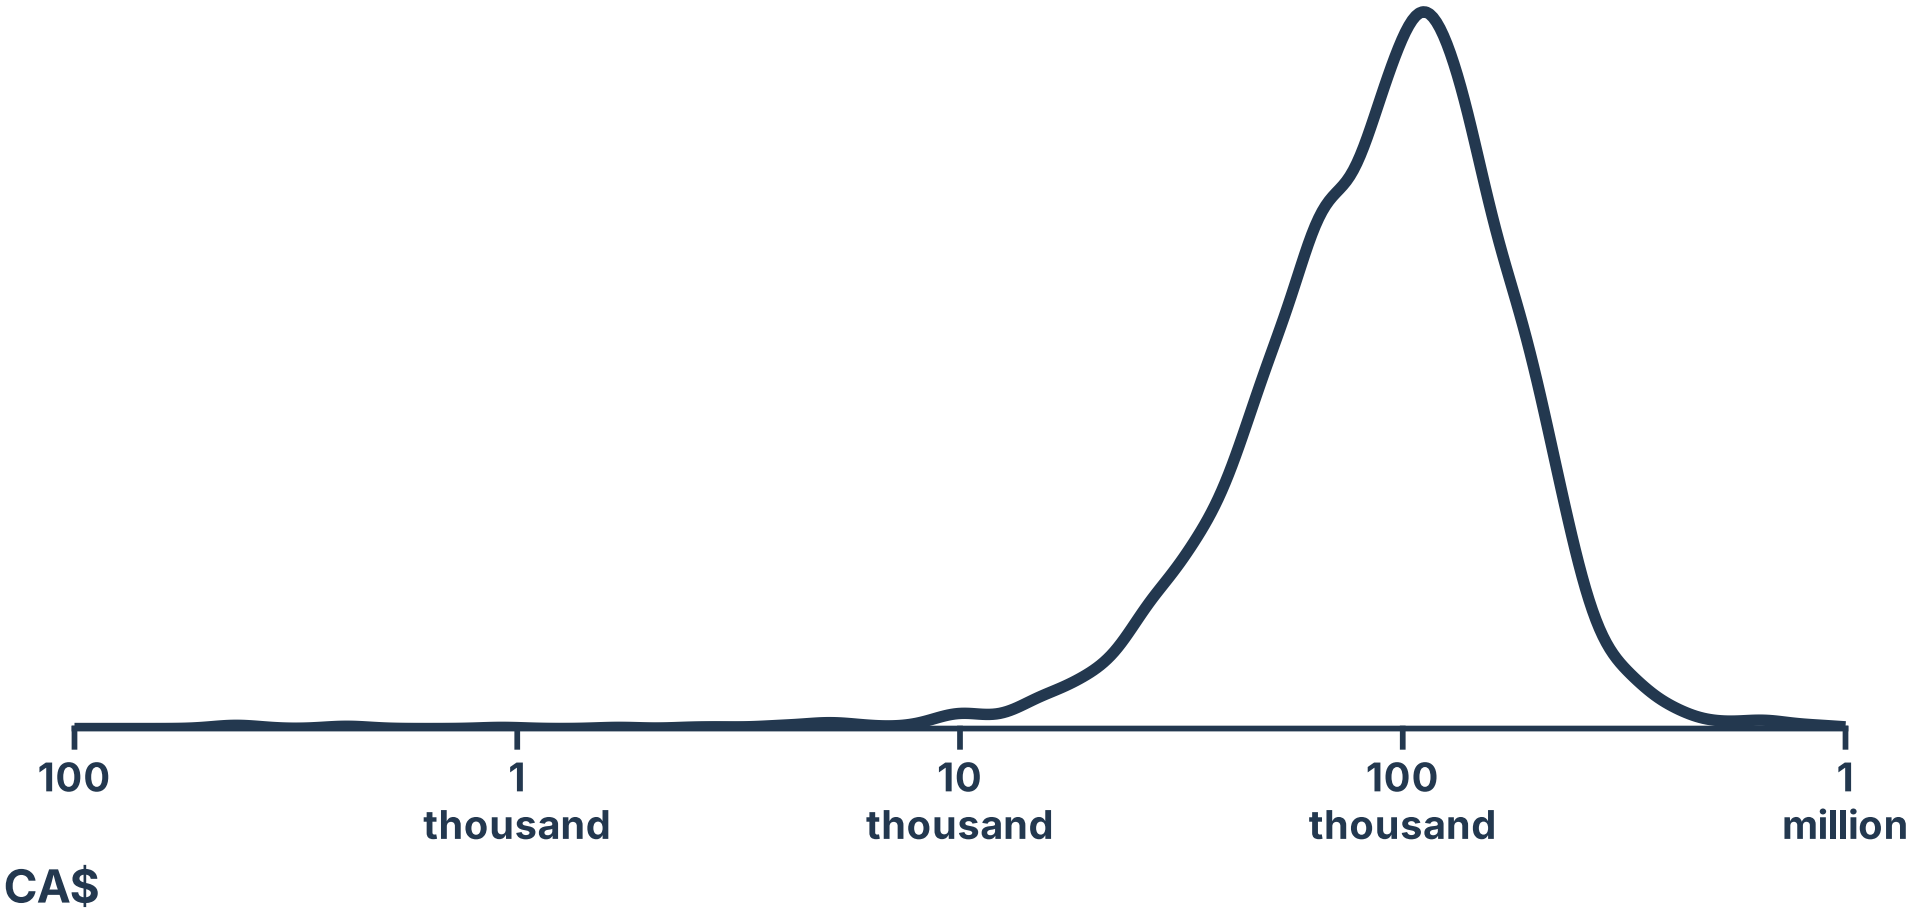

Annual family income in Canada, 2021

(weighted subsample of employed; n=3195)

Data source:
Canadian Income Survey: Public Use Microdata

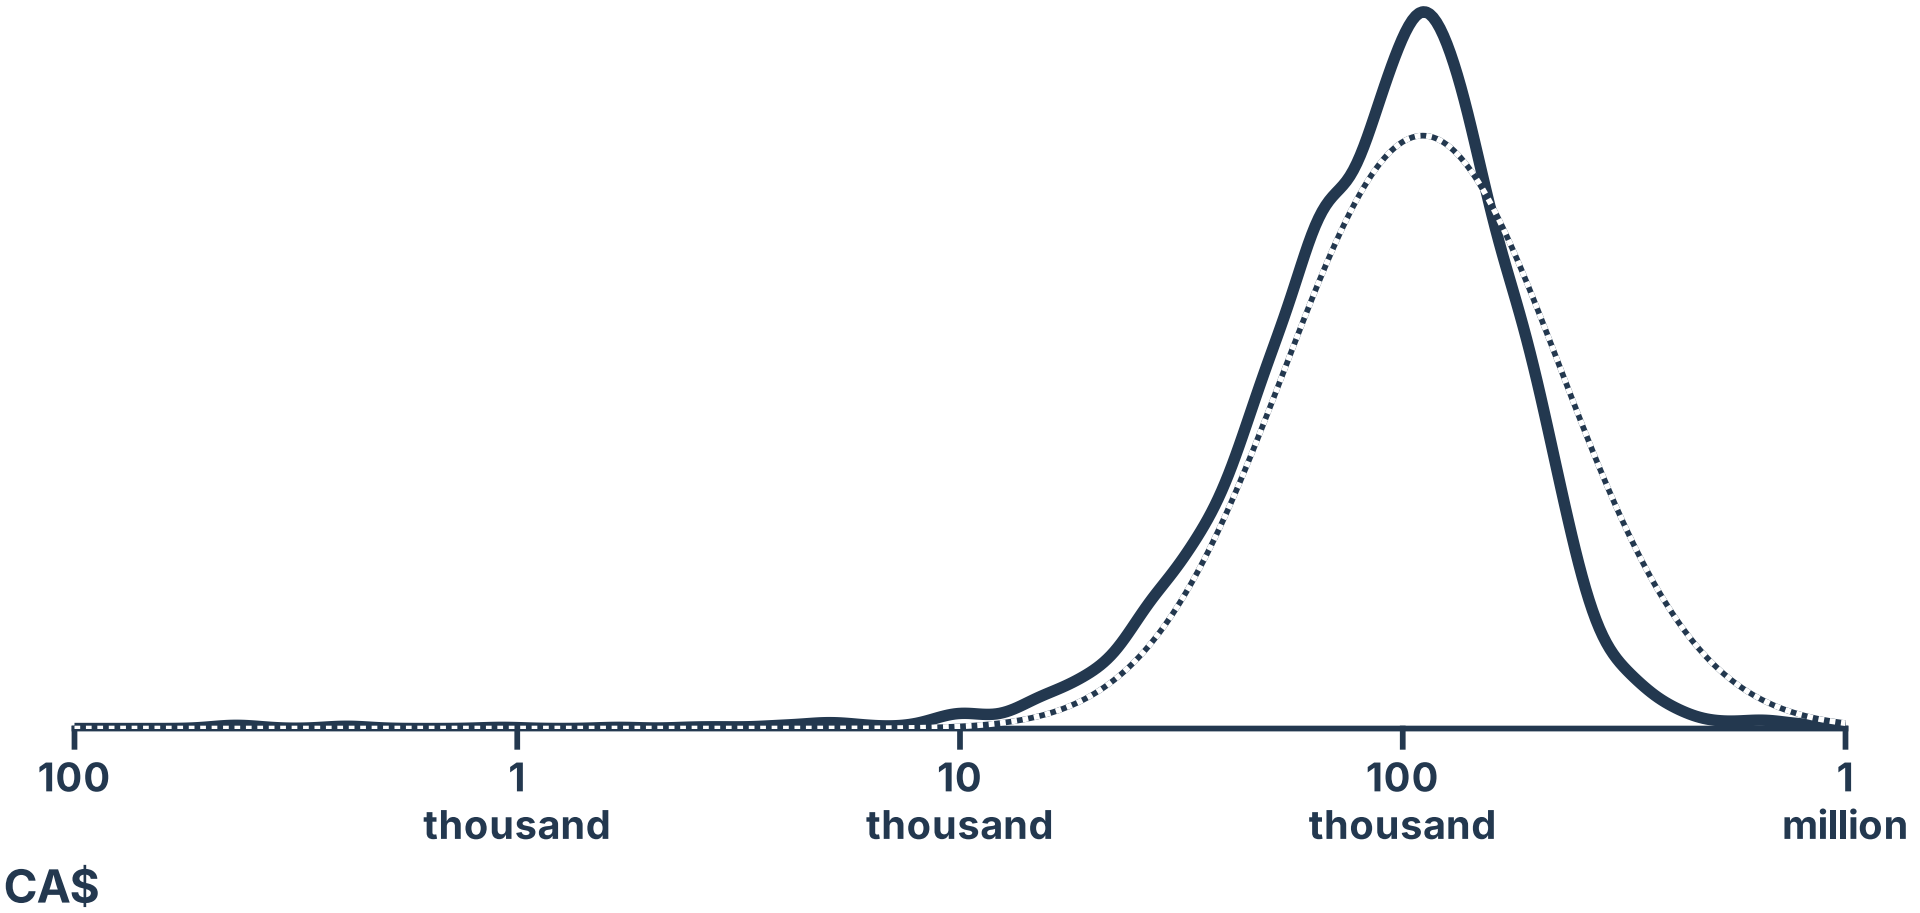

Plan: Model *log* income as a normally distributed random variable

A TWO-PARAMETER MODEL

Model of normally distributed log income:

$$y_i \sim \text{Norm}(\mu, \sigma)$$

Two Parameters:
μ determines 'location'
σ determines 'width'

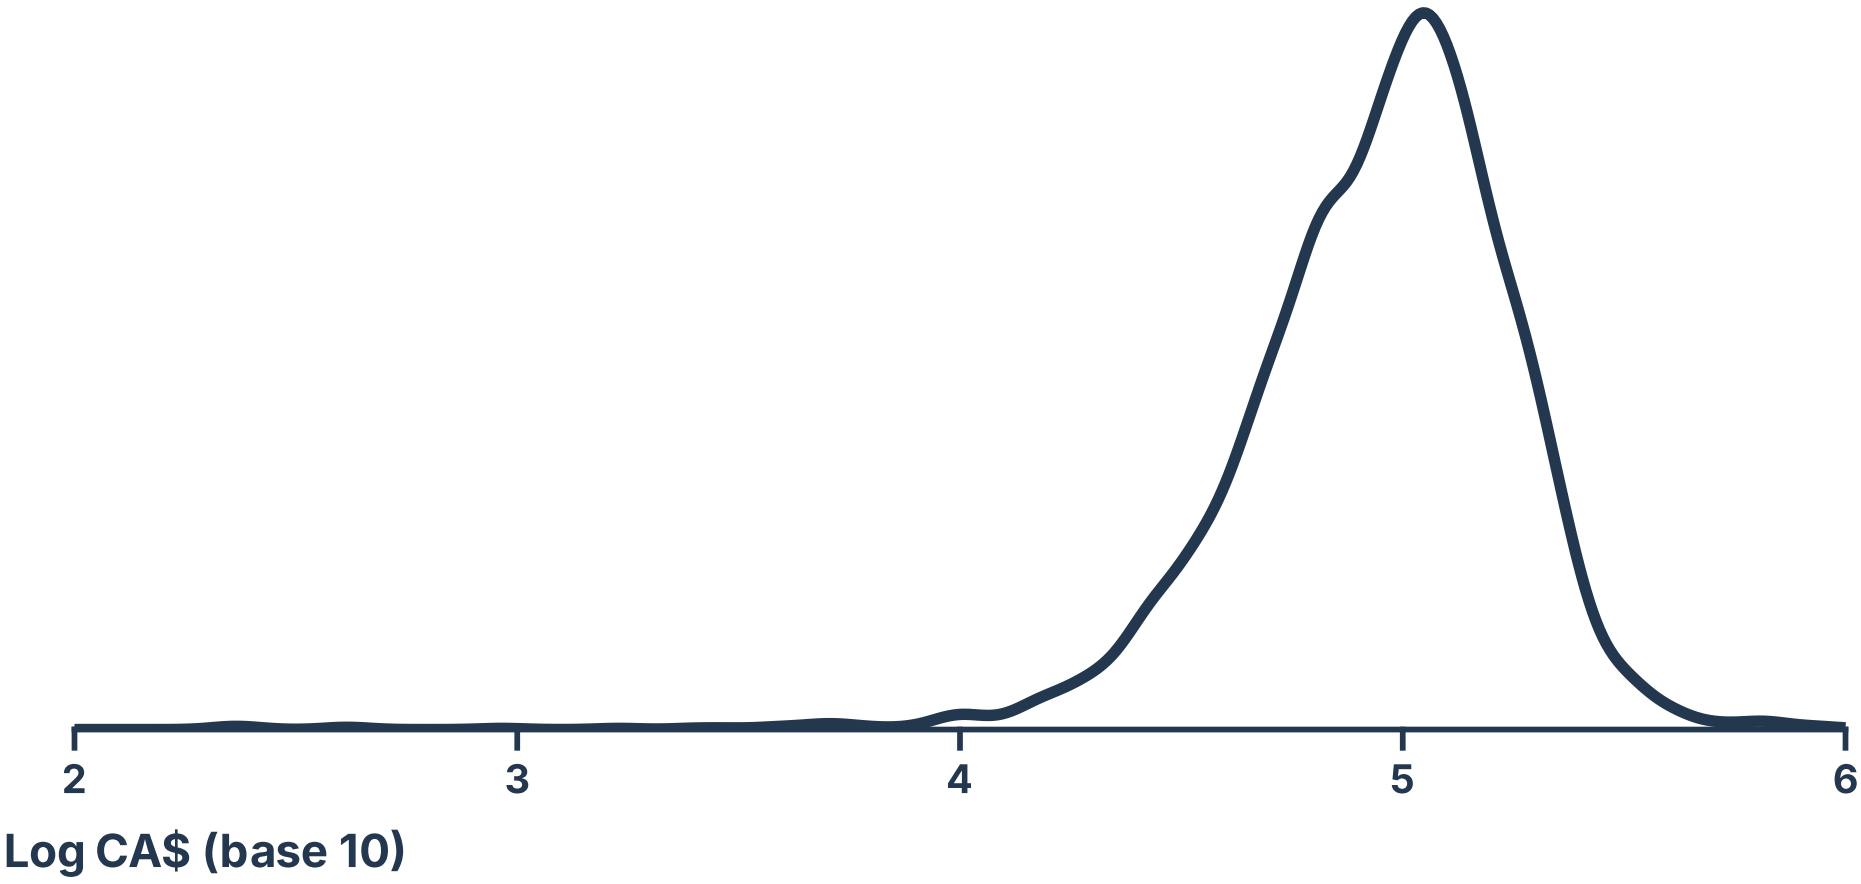

A TWO-PARAMETER MODEL

Model of normally distributed log income:

$$y_i \sim \text{Norm}(\mu, \sigma)$$

$$\mu \sim ???$$

A TWO-PARAMETER MODEL

Model of normally distributed log income:

$$y_i \sim \text{Norm}(\mu, \sigma)$$
$$\mu \sim \text{Norm}(0, 30)$$

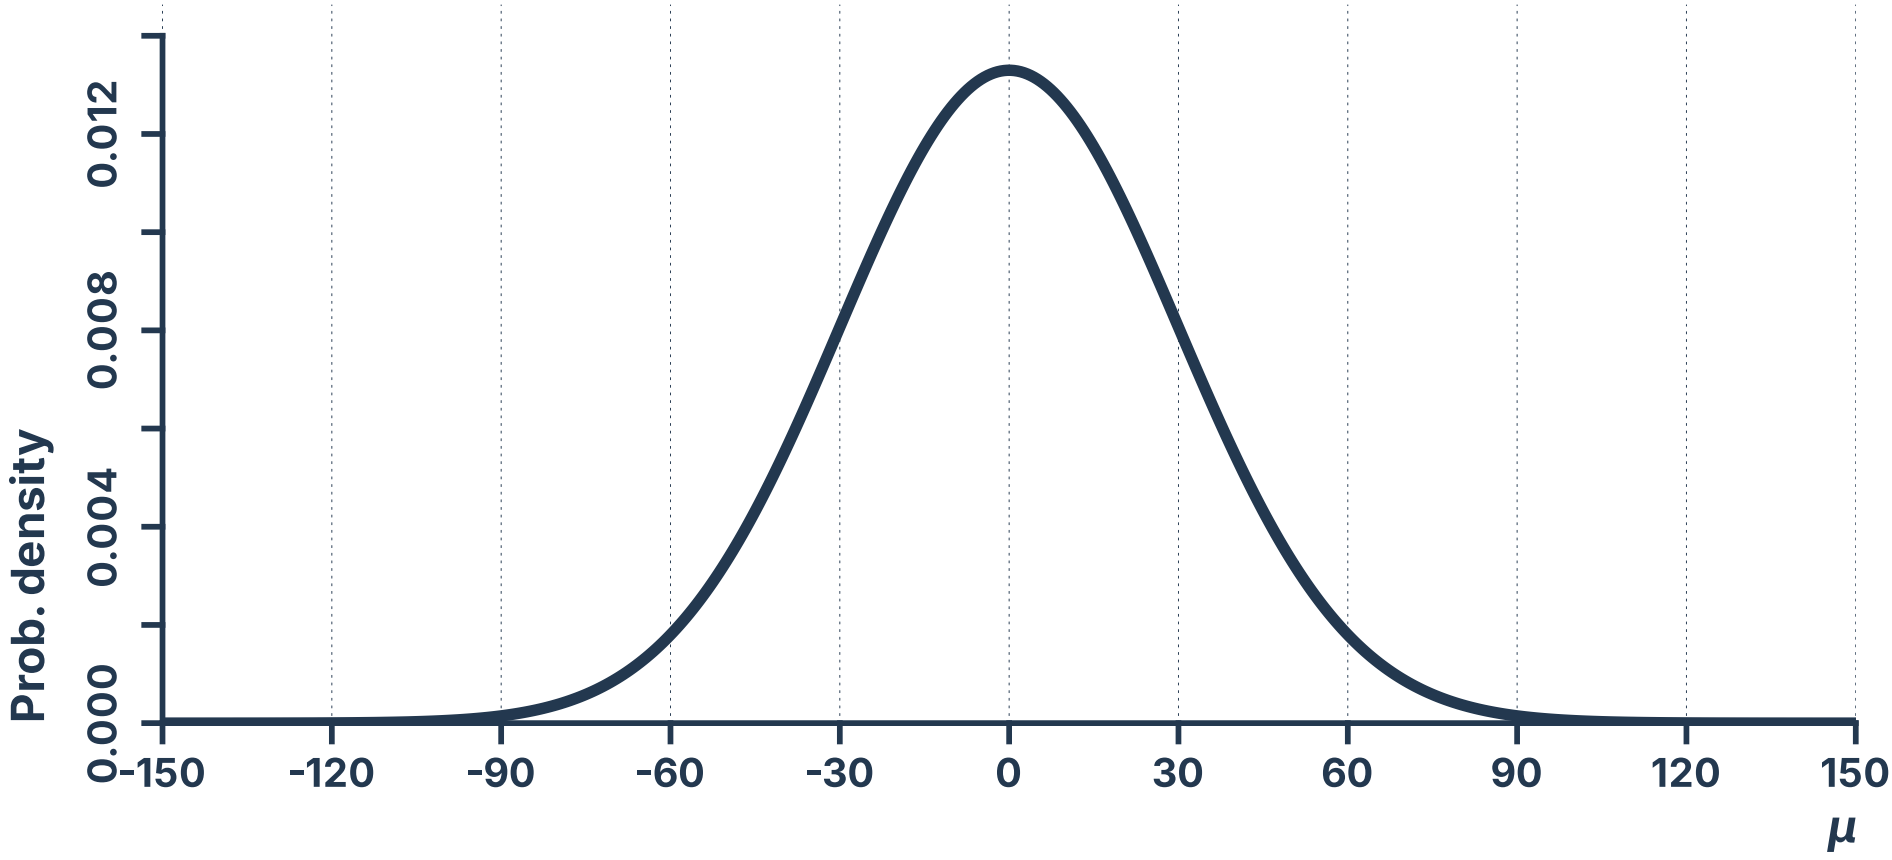

A TWO-PARAMETER MODEL

Model of normally distributed log income:

$$y_i \sim \text{Norm}(\mu, \sigma)$$
$$\mu \sim \text{Norm}(0, 30)$$
$$\sigma \sim \text{Unif}(0, 50)$$

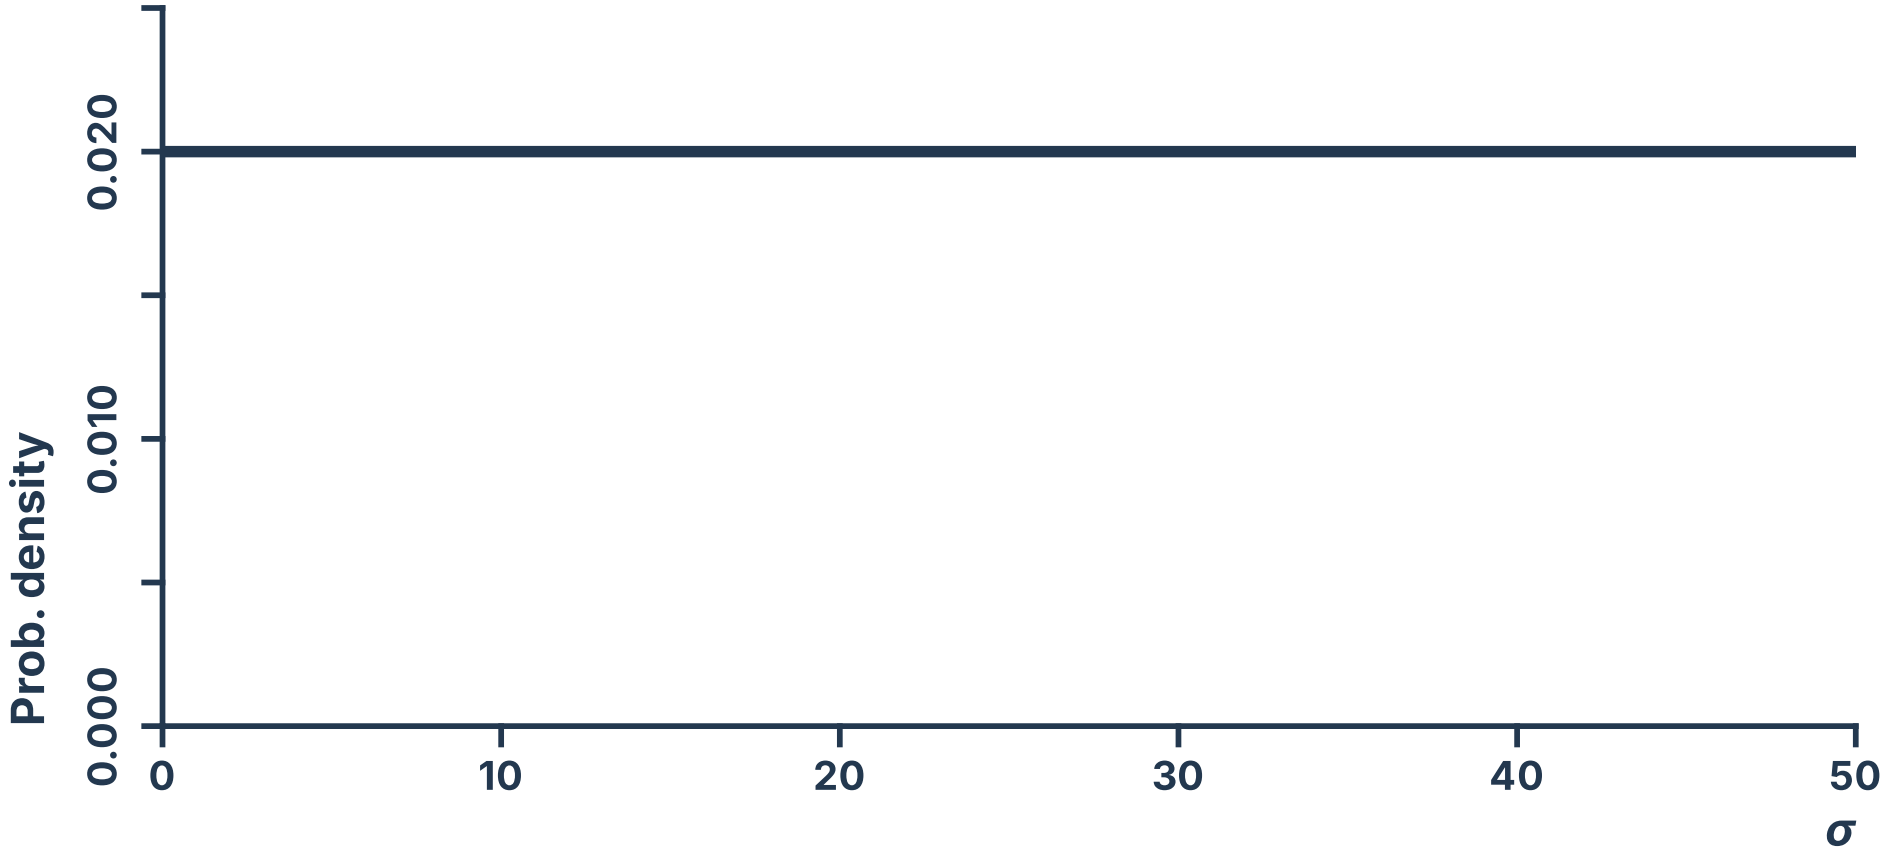

A TWO-PARAMETER MODEL

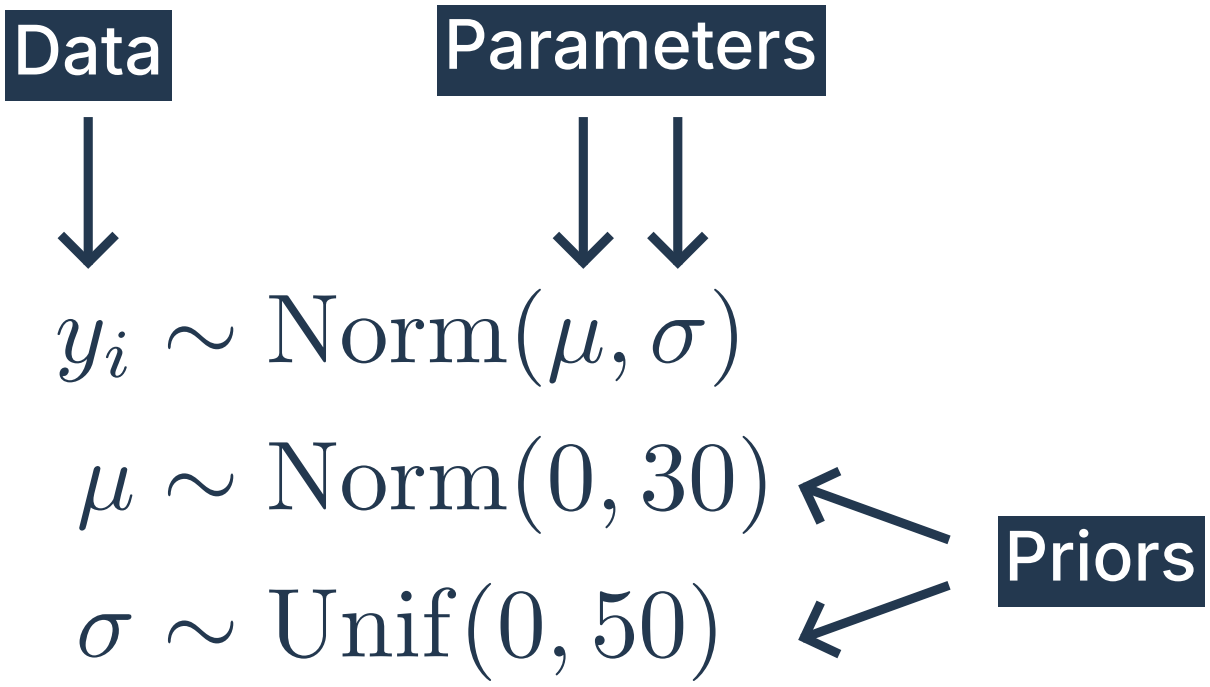

MULTI-PARAMETER POSTERIORS

Pr(μ) prior

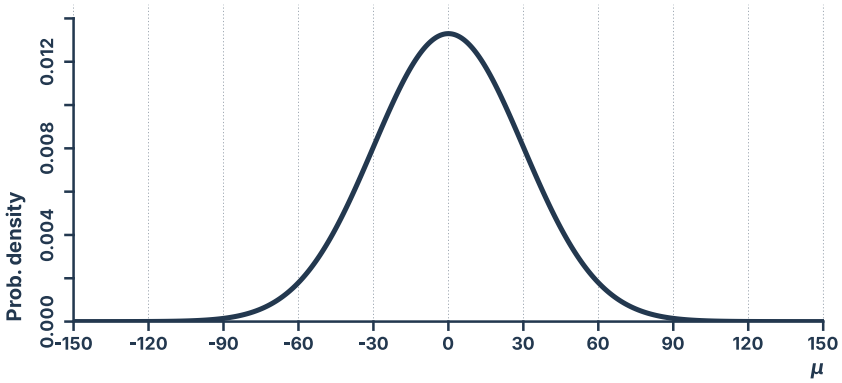

Pr(μ | Y) posterior

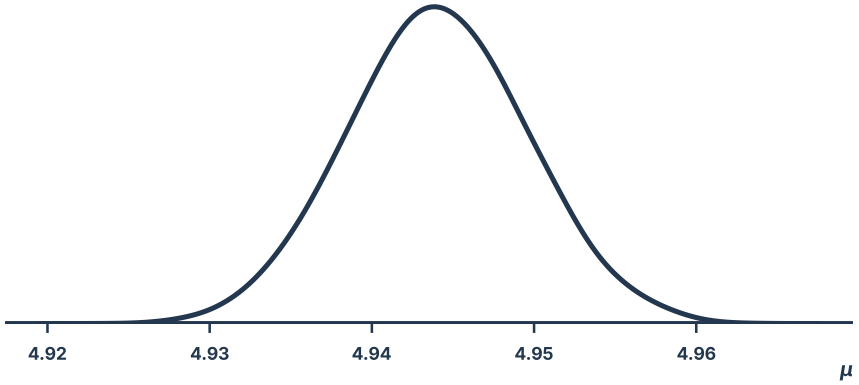

Pr(σ) prior

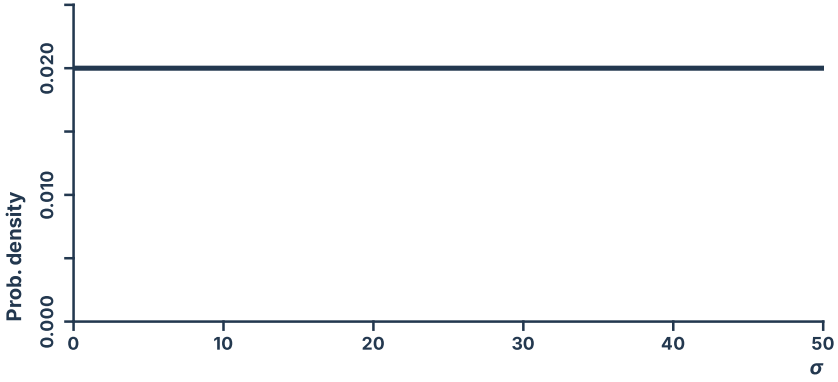

Pr(σ | Y) posterior

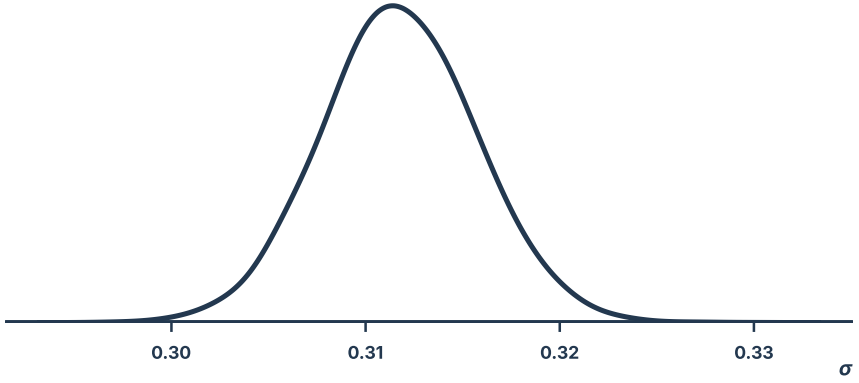

MULTI-PARAMETER POSTERIORS

Pr(μ,σ | Y) posterior

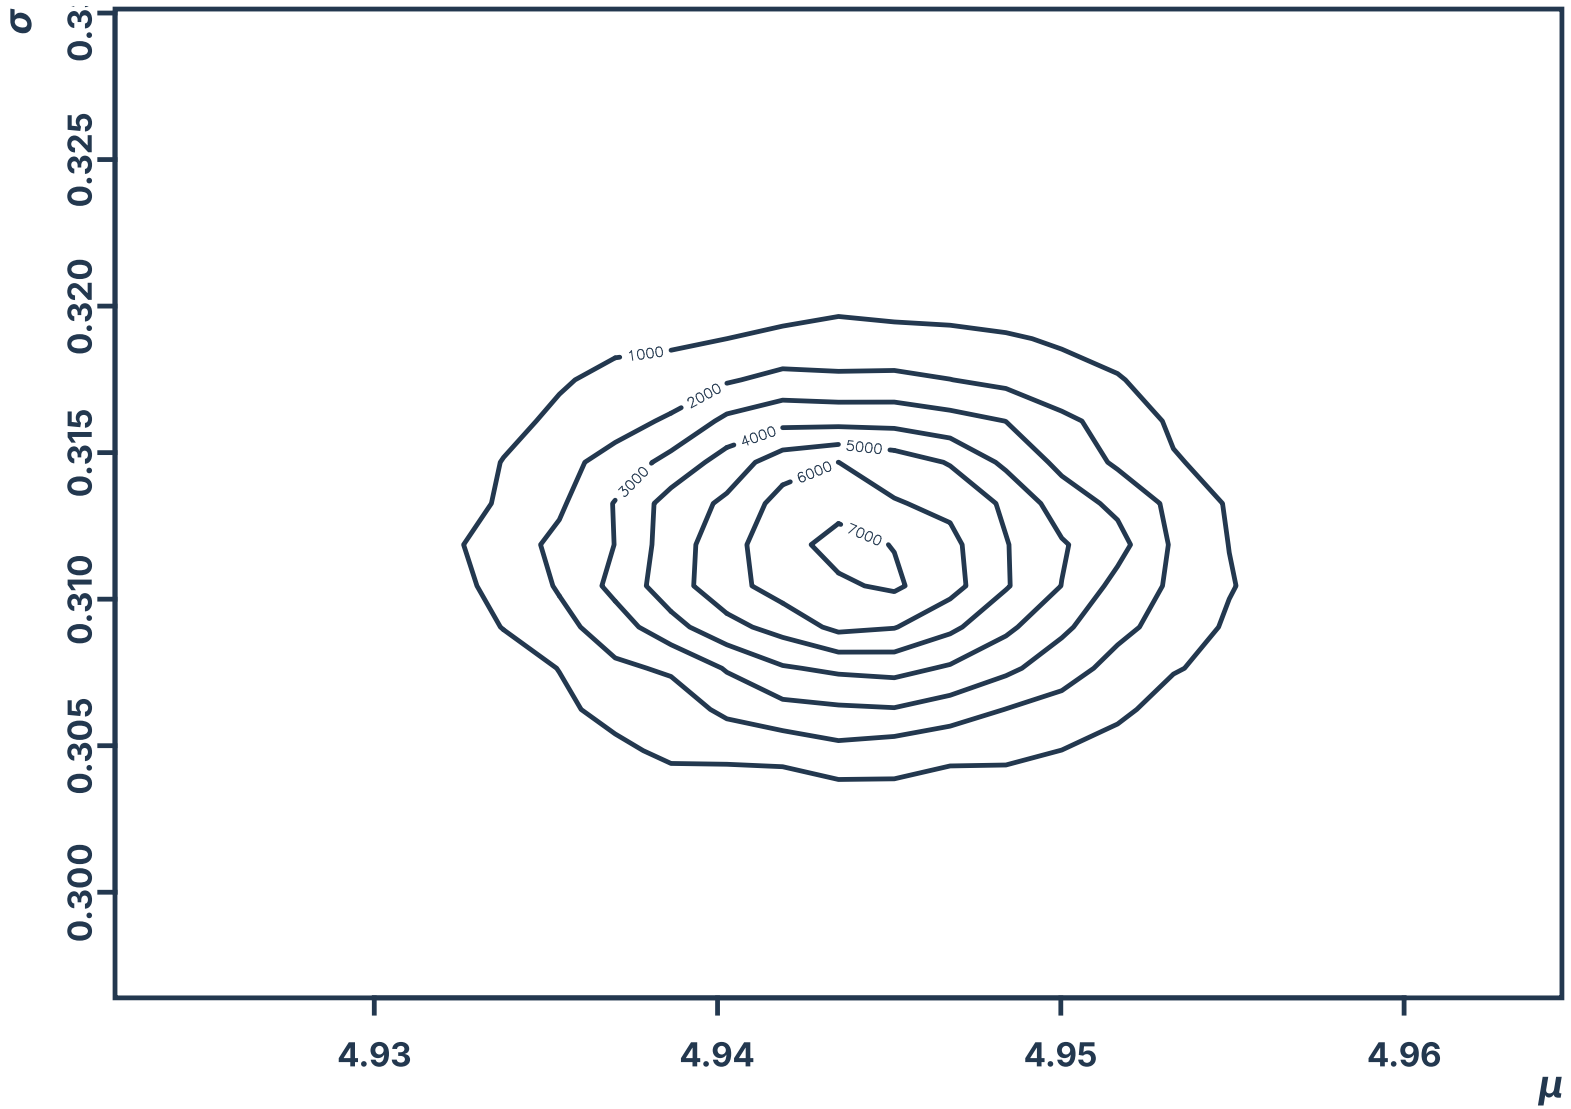

Maximum a posteriori estimation

(MAP/QUAP)

MAP ESTIMATION

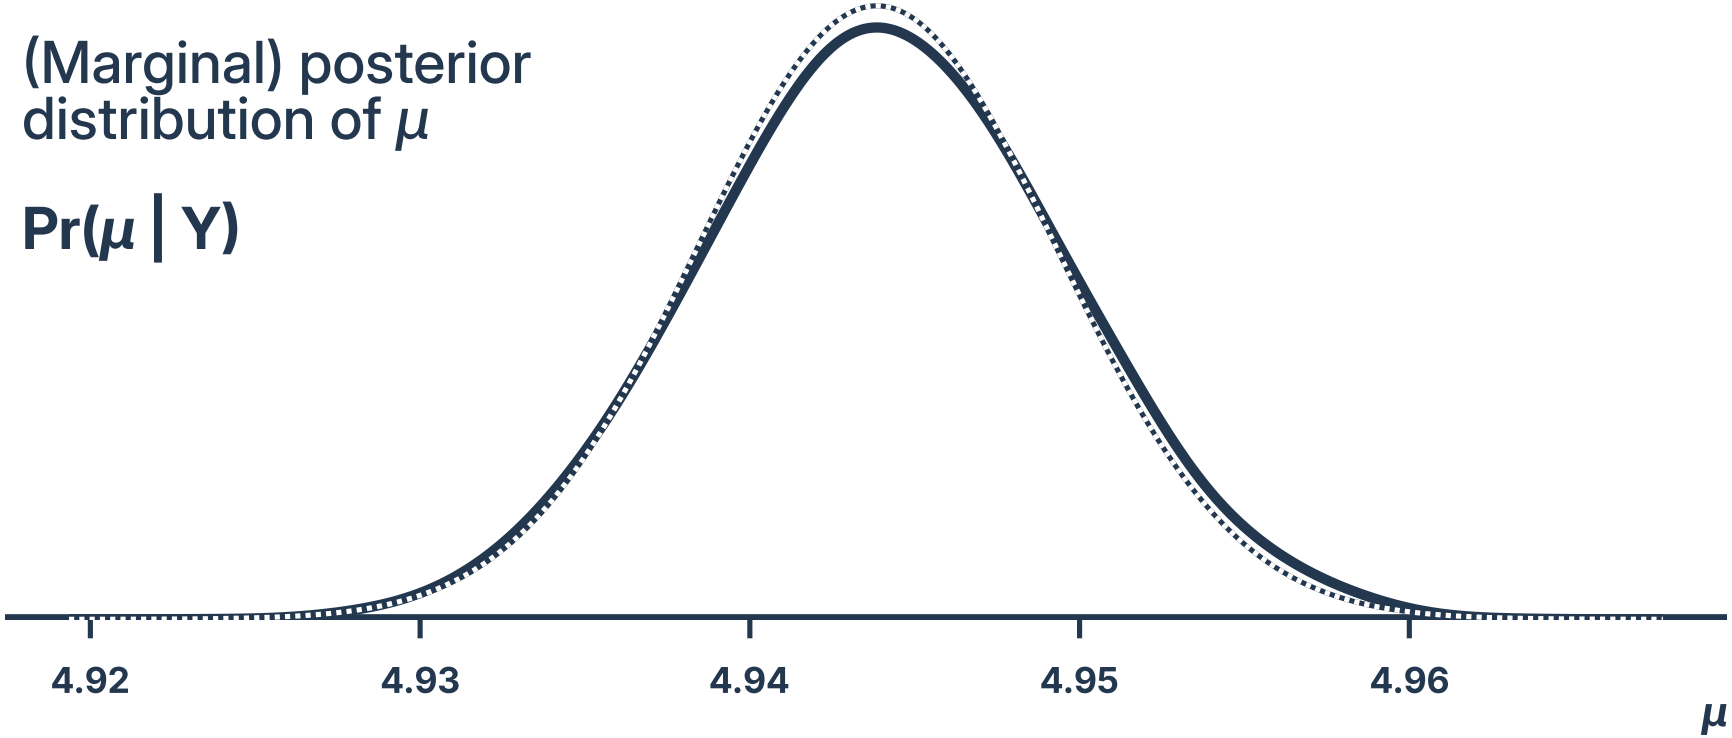

Maximum a posteriori (quadratic approximation)

1. Find the 'peak' of the posterior and call it m_μ
2. Estimate the 'spread' of the posterior and call it s_μ
3. Approximate the posterior with Norm(m_μ, s_μ)

NORMALS ALL THE WAY DOWN

So many normal distributions!
This can get confusing!

Pr(Y | μ)

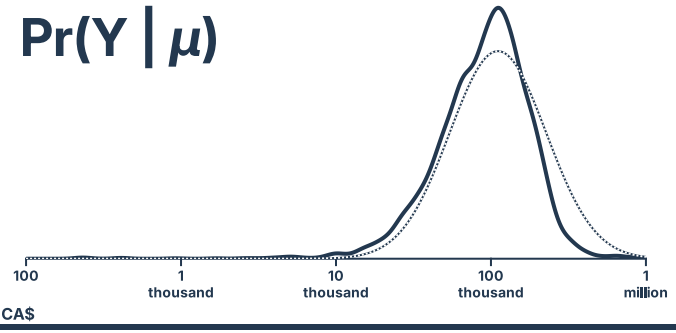

- ***First*** we modeled income using a normal distribution with unobserved parameters μ and σ:

$$y_i \sim \text{Norm}(\mu, \sigma)$$

Pr(μ)

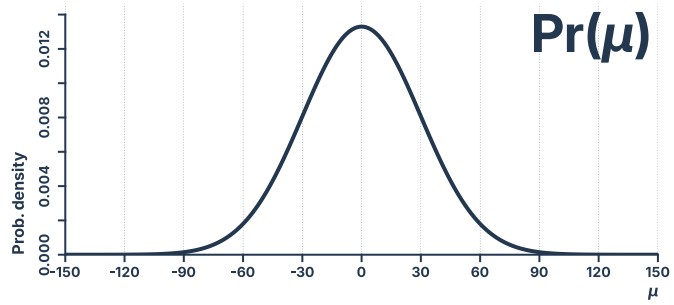

- ***Then*** we gave σ a normal prior:

$$\mu \sim \text{Norm}(0, 30)$$

Pr(μ | Y)

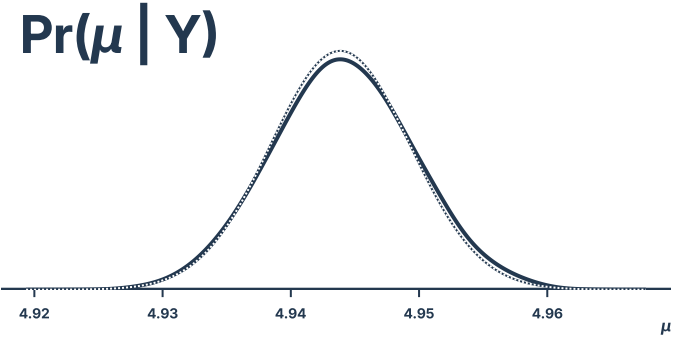

- And ***then*** we used a normal distribution to approximate the *posterior* distributions of μ and σ:

$$\mu|Y \sim \text{Norm}(m_\mu, s_\mu)$$

$$\sigma|Y \sim \text{Norm}(m_\sigma, s_\sigma)$$

Image credit

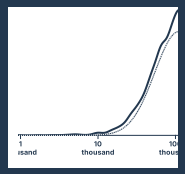

Figures by Peter McMahan (source code)

Image copyright Disney

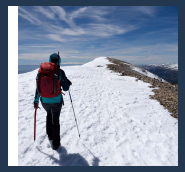

Photo by Paulo Etxeberria