Sampling distributions

Topic 6 · Thursday, January 21 and Tuesday, January 26

EDUC 15: Statistics for Education Research
UCI School of Education

Where are we?

Description

Prediction

Explanation

Summarize data

Visualize data

We are here, heading right



Peter the Anteater at a UC Irvine basketball game.

The question we keep asking

If what I think about the population is true, how weird is my sample?

So far, a sample was one person. We asked how weird that one score was

A real sample has many people. So we compute its mean, and ask how weird that mean is

Then, how weird must it be before I change my mind? That's Topic 7

Plan for the rest of today

The distribution of sample means

How weird is one mean?

On Tuesday: standard errors, how to estimate them, and how to use them

The distribution of sample means

Our blueprint: three panels, one story

Panel 1: The population

Everyone we care about. Its μ and σ are fixed, and usually hidden from us

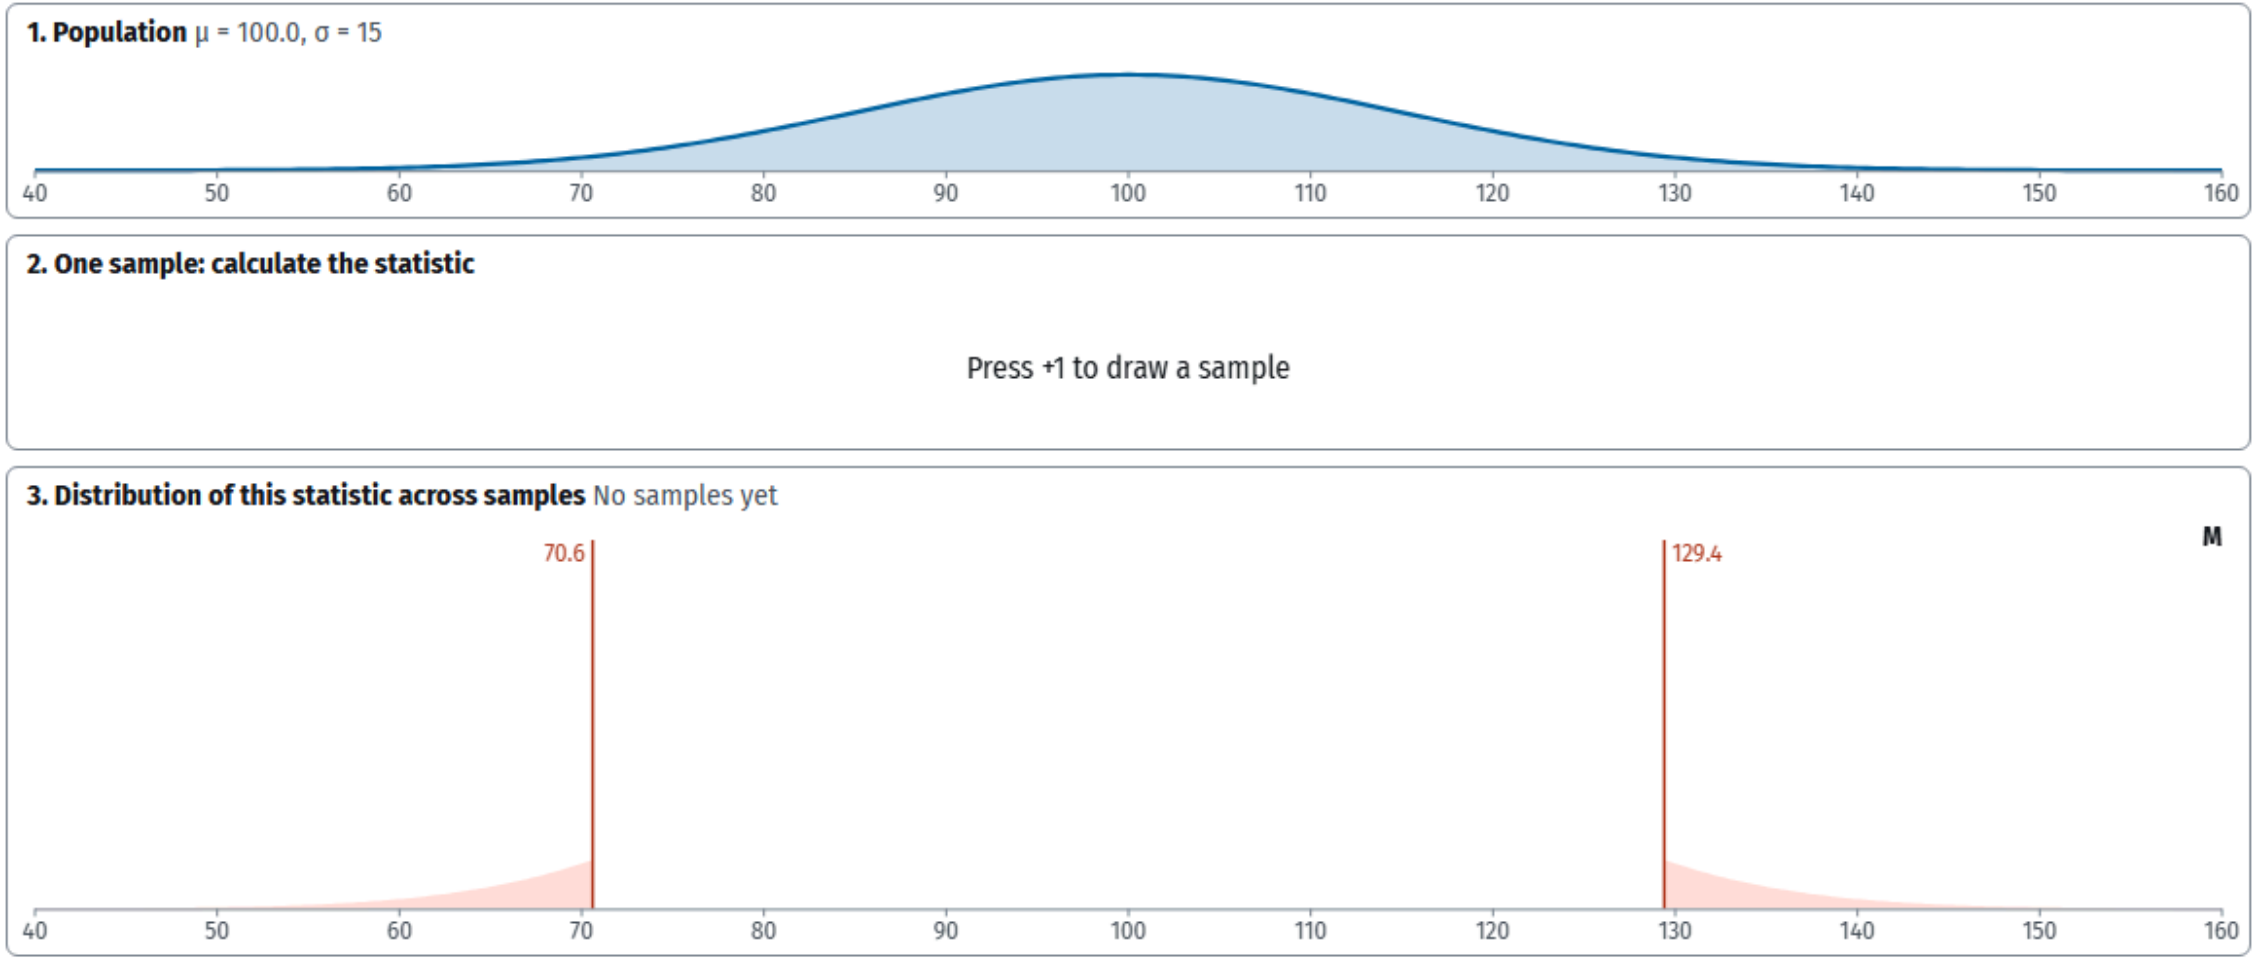
Panel 2: One sample

We draw n people and compute one statistic, such as their mean M

Panel 3: That statistic, across samples

We repeat, and watch how the statistics stack up

Start small: samples of one person



Simulator after Kristoffer Magnusson, rpsychologist.com.

That gap is called sampling error

In points

$M - \mu = 106.2 - 100 = \mathbf{6.2}$

In SDs: a z-score

$(106.2 - 100) / 15 = \mathbf{0.41}$

Sampling error is the gap between a statistic and the parameter it estimates

Nobody made a mistake. A random sample rarely matches its population exactly

Predict it: Fifty more samples of one

We have one dot in panel 3, at 106.2. We draw 50 more samples, then 500 more. Each gives one mean, and each mean is one dot

What will the 551 dots look like?

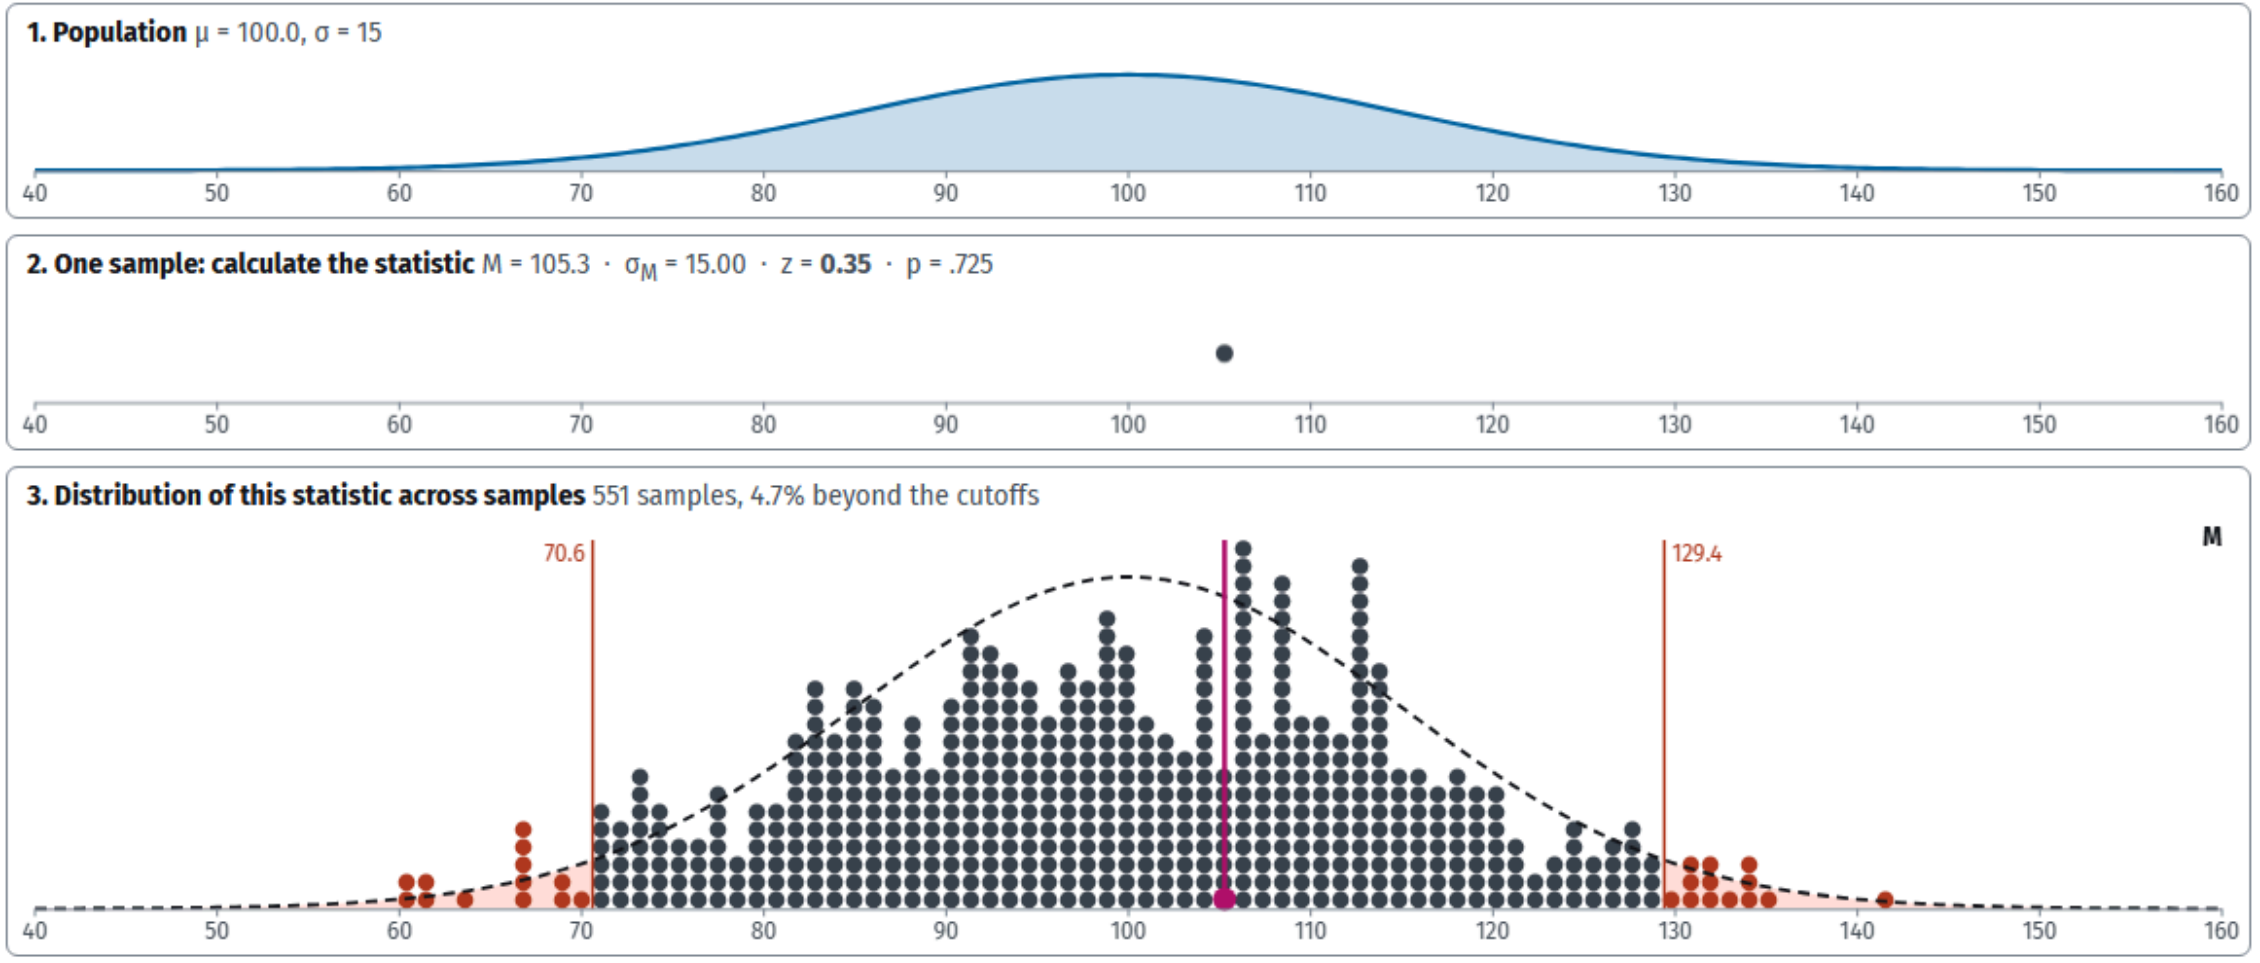
A A bell around 100, as wide as the population

B A narrow spike right at 100

C A flat line across the whole axis

D A bell around 106.2, our first sample

n = 1: the means copy the population



Simulator after Kristoffer Magnusson, rpsychologist.com.

How weird is one mean?

That pile has a name

Sampling distribution

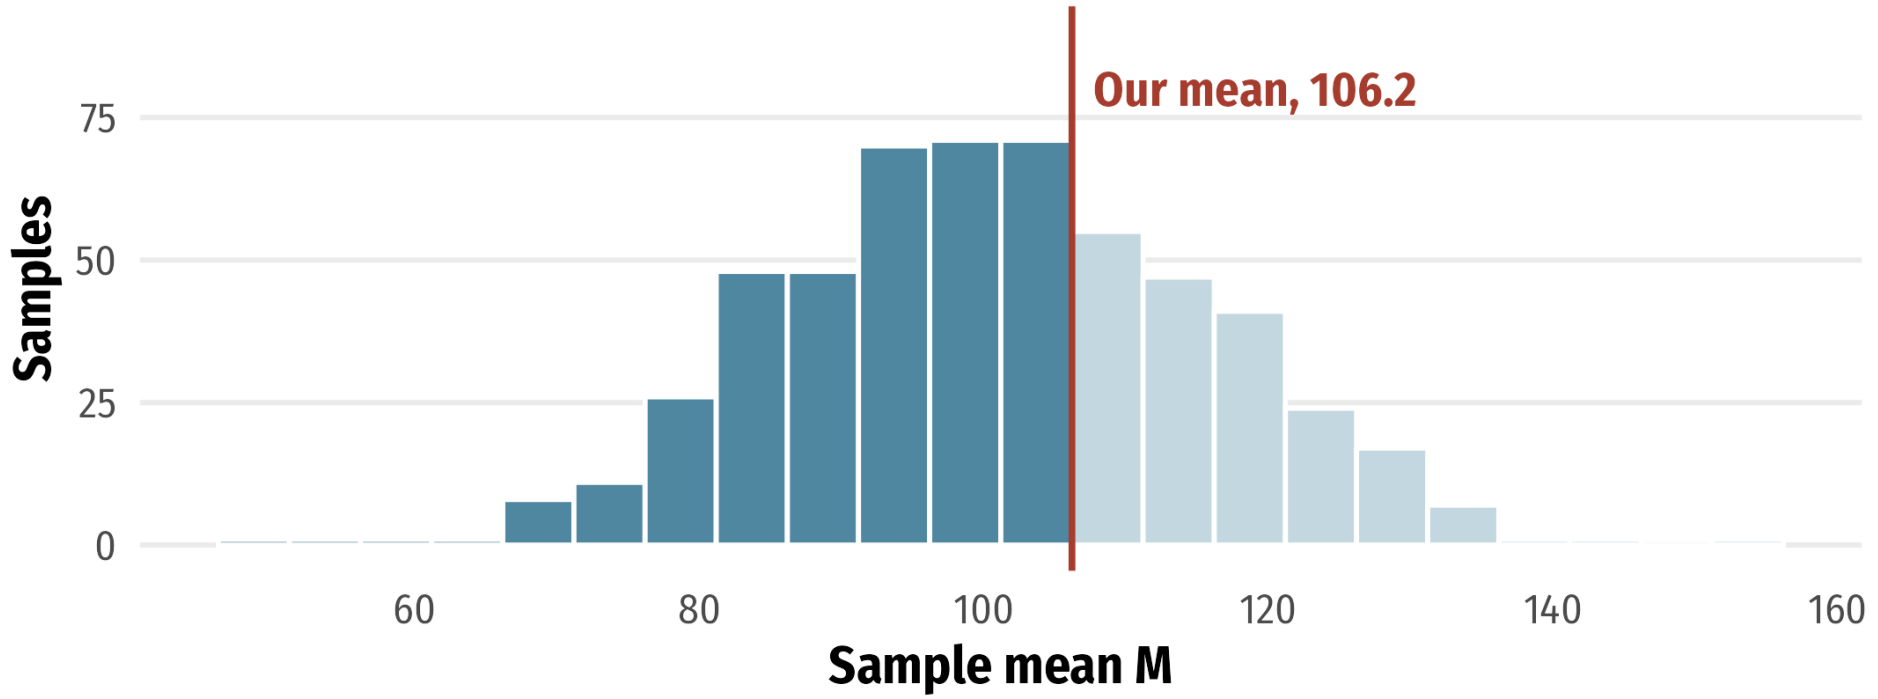
The distribution of a statistic, across all possible samples of size n

Distribution of sample means

The same thing, when the statistic is the mean M

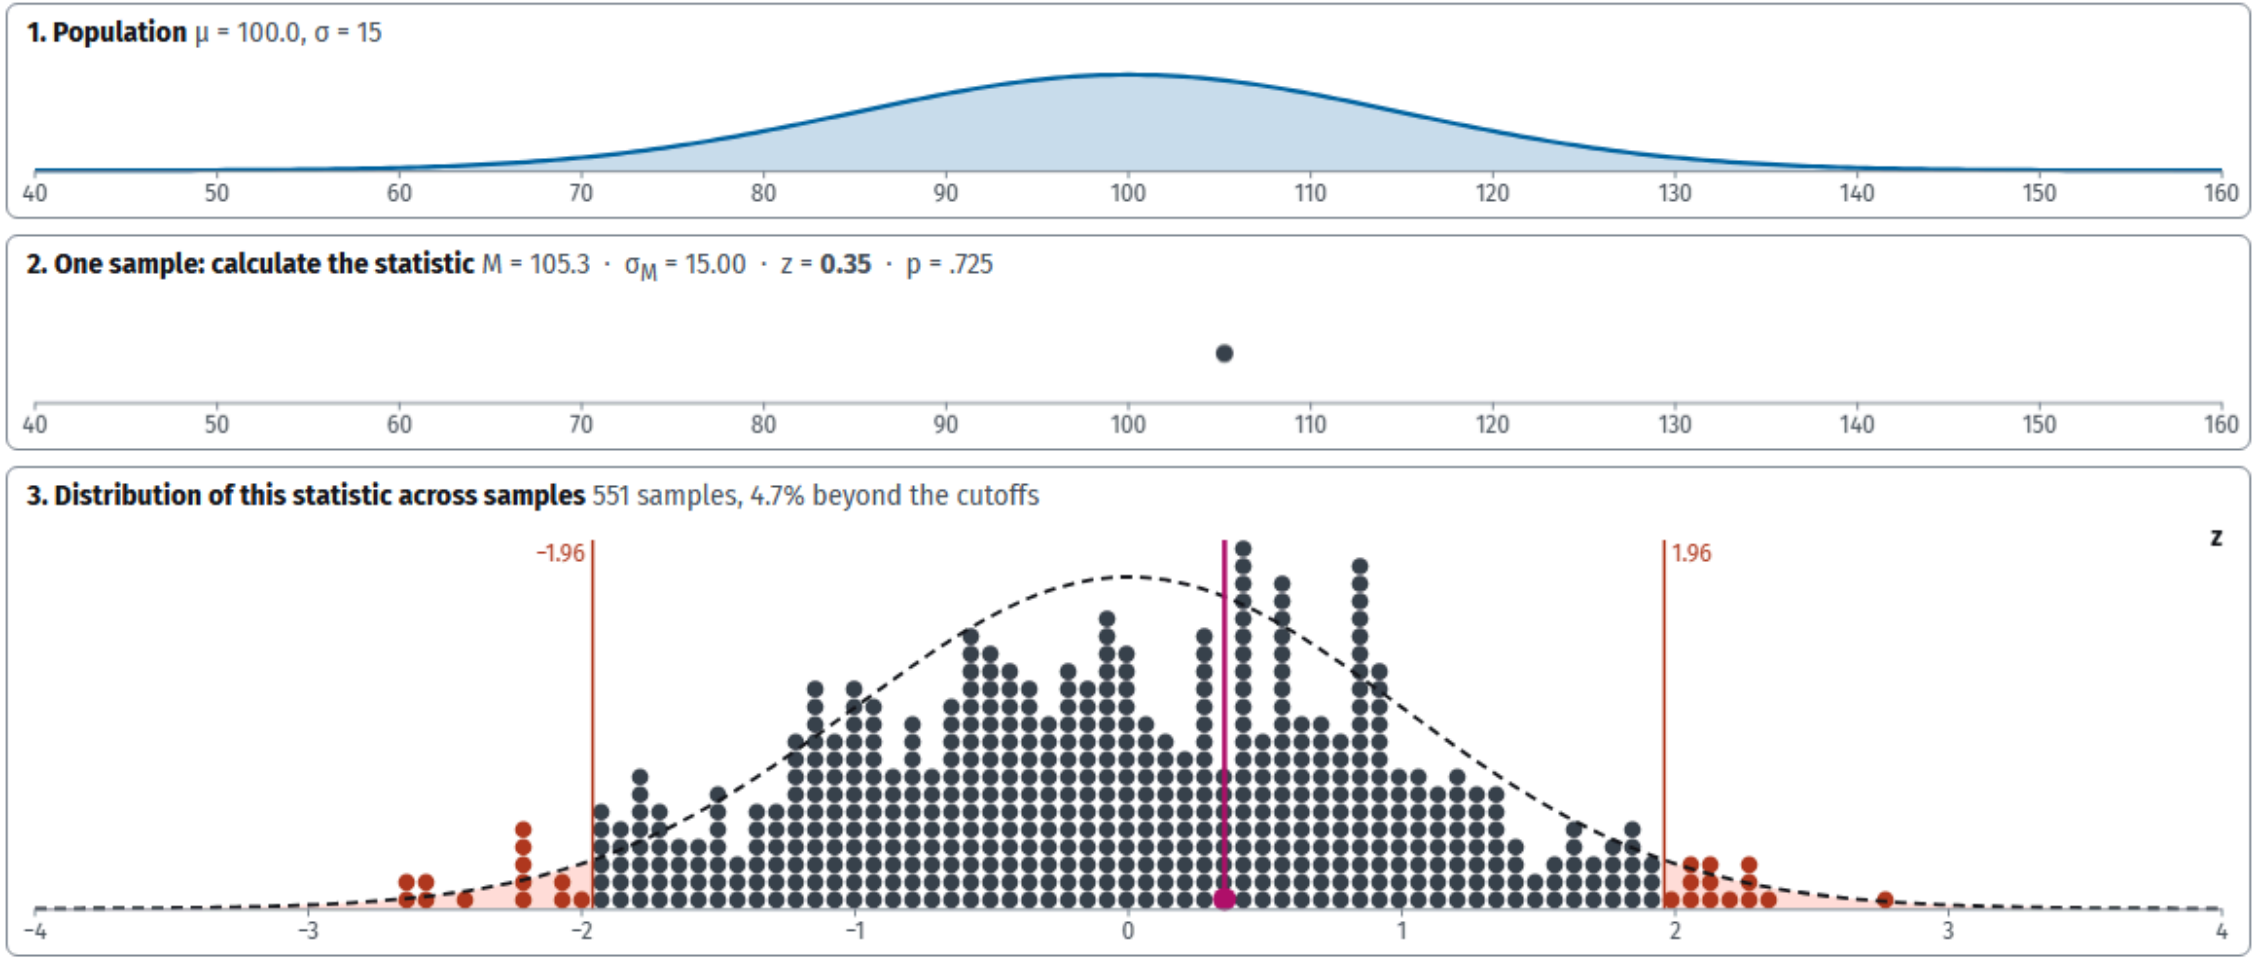
Each dot is a whole sample, not a person

Count the dots to find a probability



357 of 551 means are below 106.2. So p(M < 106.2) is about 0.65

Turn every mean into a z-score



Simulator after Kristoffer Magnusson, rpsychologist.com.

So we can skip the counting

Any normal distribution in z-scores is the standard normal, so the unit normal table applies

$$z = \frac{M - \mu}{\sigma} = \frac{106.2 - 100}{15} = 0.41$$

Table: the body at z = 0.41 is 0.6591

Counting gave 0.65

A mean like ours isn't weird. About 66% of samples fall below it

Predict it: A lopsided population

Now the population is skewed right. We draw samples of 30 and compute each mean

What shape will the means make in panel 3?

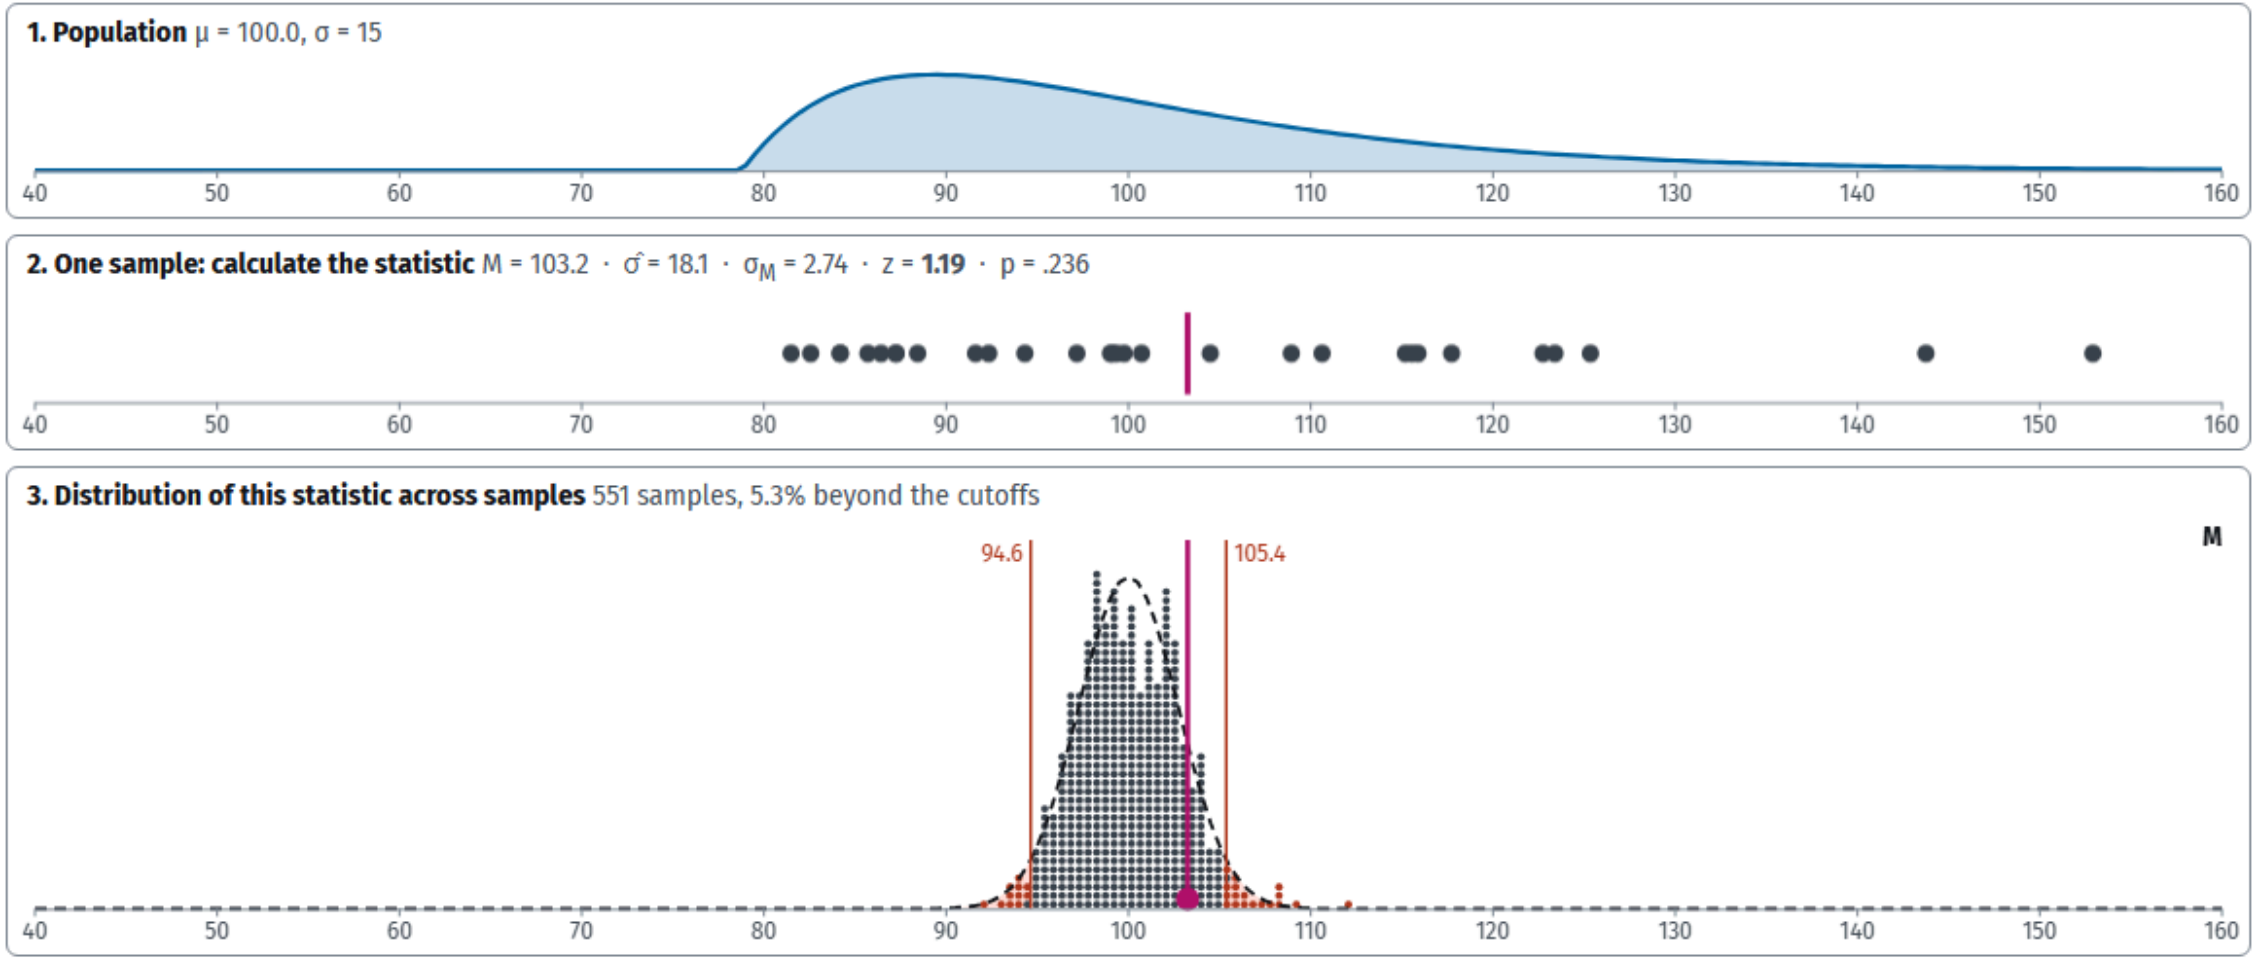
A Skewed right, like the population

B A bell around the population's mean

C Skewed left, the opposite way

D Flat

Skewed population, bell-shaped means



Simulator after Kristoffer Magnusson, rpsychologist.com.

The central limit theorem

With large samples, the distribution of sample means is close to normal, whatever the population's shape

If the population is normal, the means are normal at any n

If it isn't, n of about 30 is usually enough

So with a large n, the unit normal table works for means from any population

Before Tuesday

Friday, January 22: Lab 3 is a quiz. Bring a calculator

Sunday, January 24 at 6 PM: Mastery and Problem Set for Chapter 6

Read sections 7-3 to 7-5 for Tuesday

Session 2: Tuesday, January 26

Thursday in one line: means of samples pile up into a distribution of their own

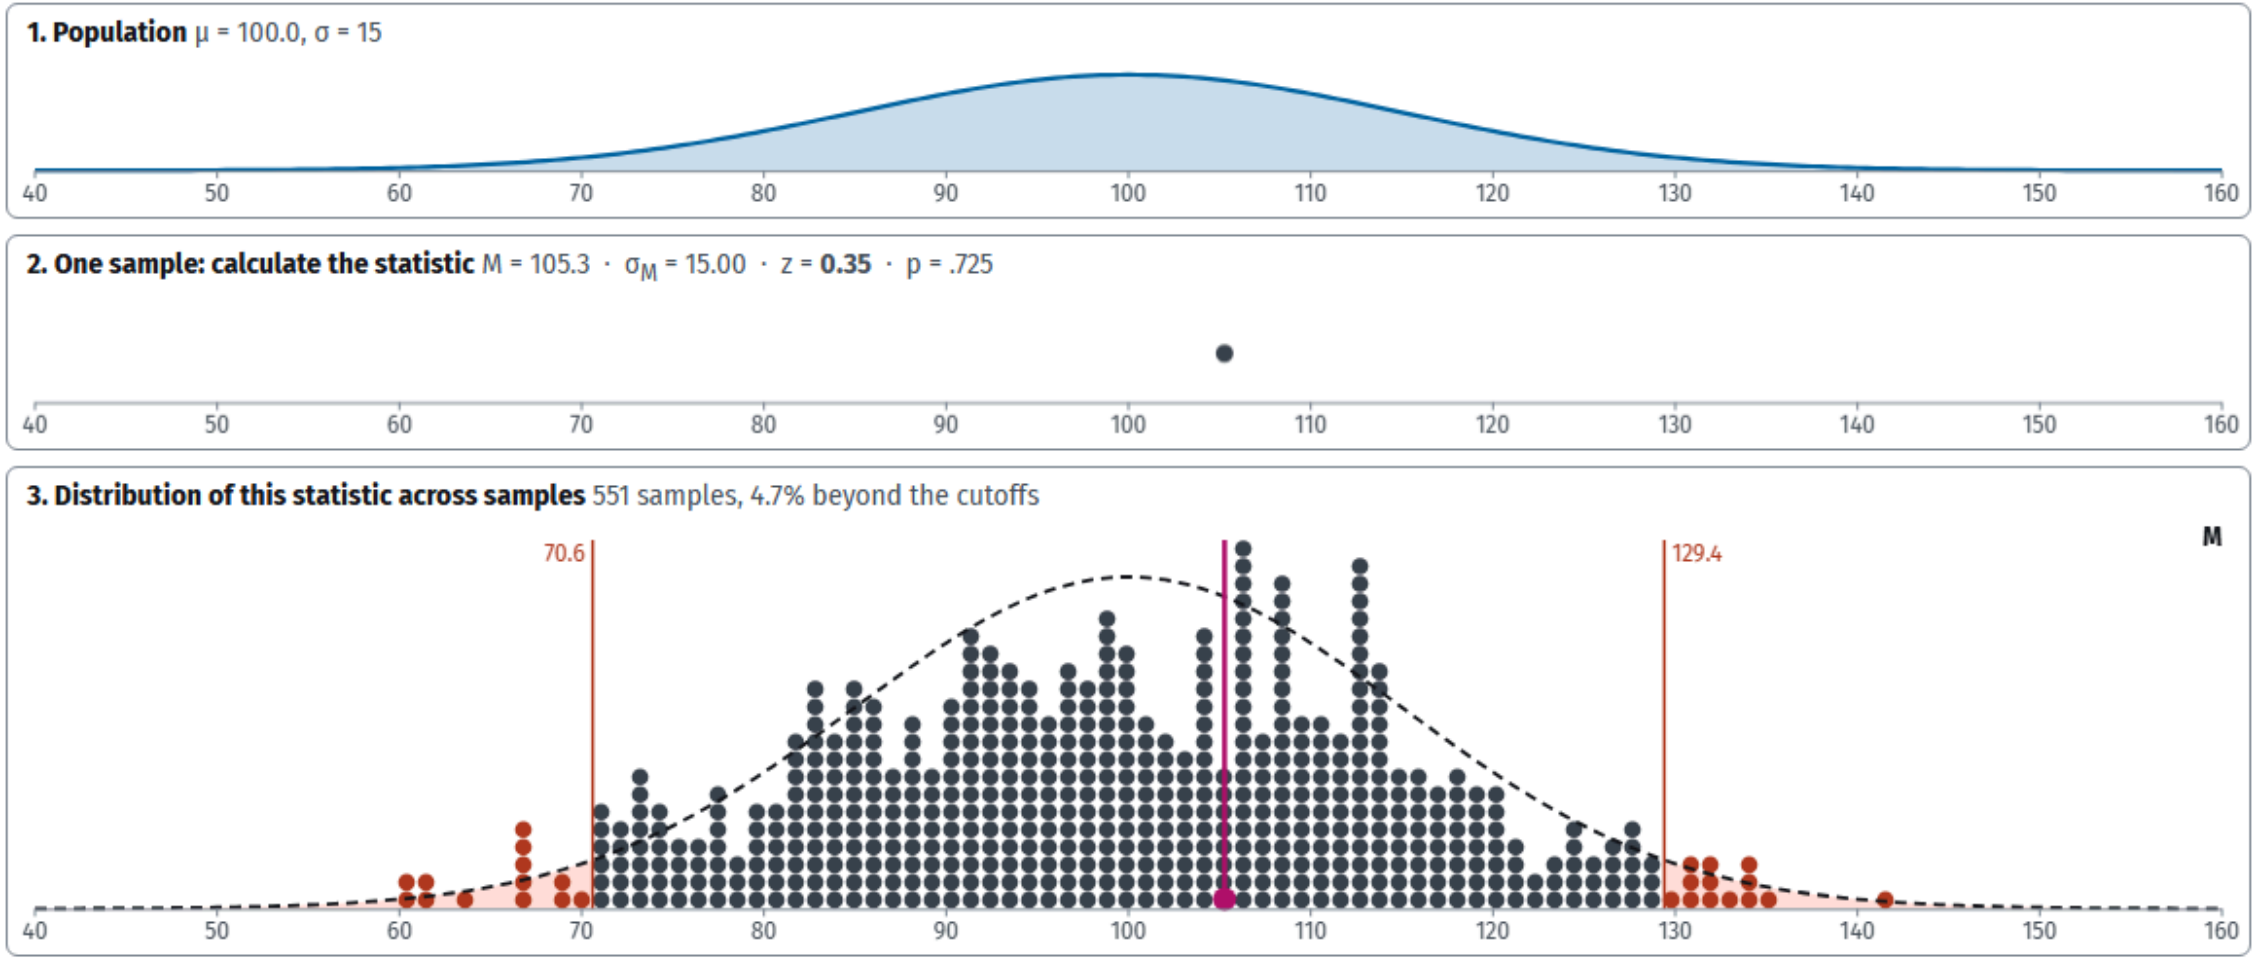
Standard errors

Estimating the standard error

Using standard errors

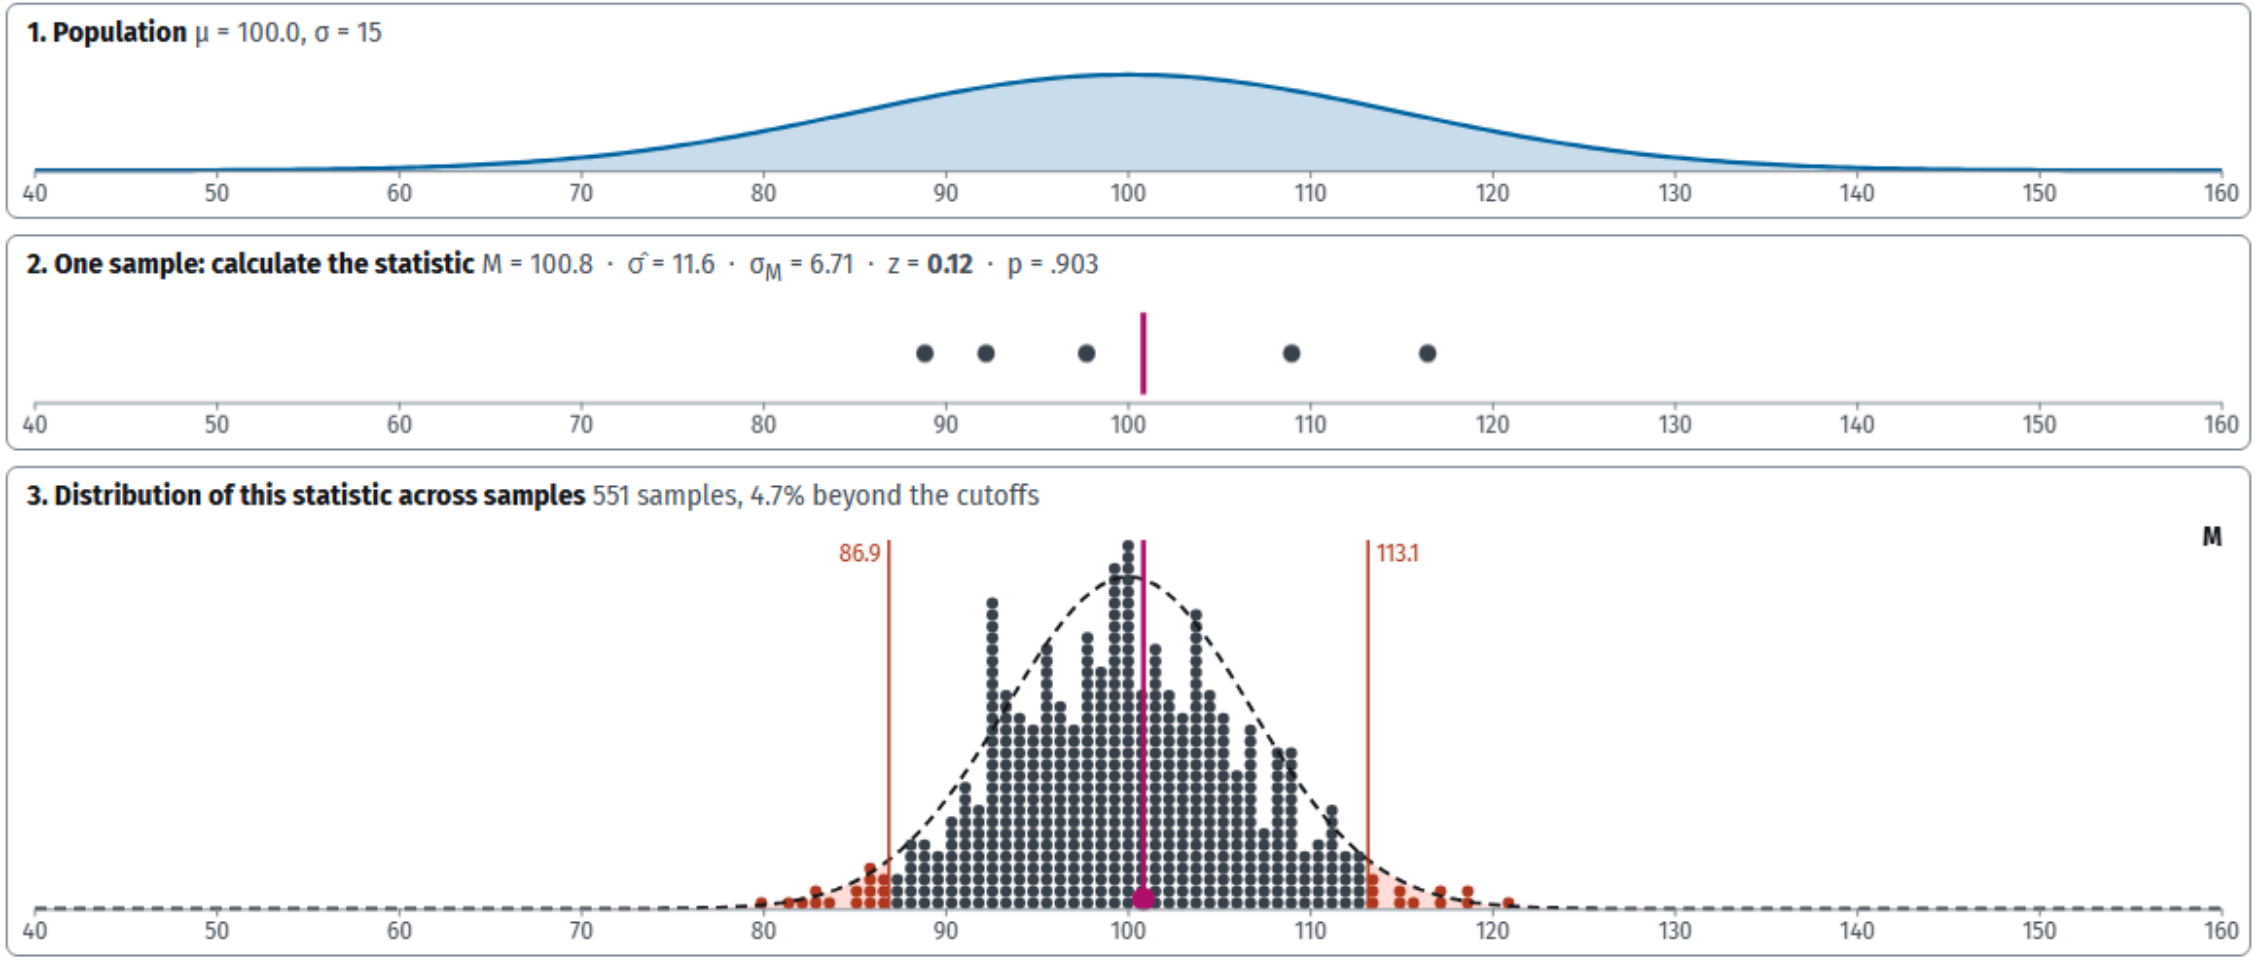
Standard errors

So far, each sample had one person



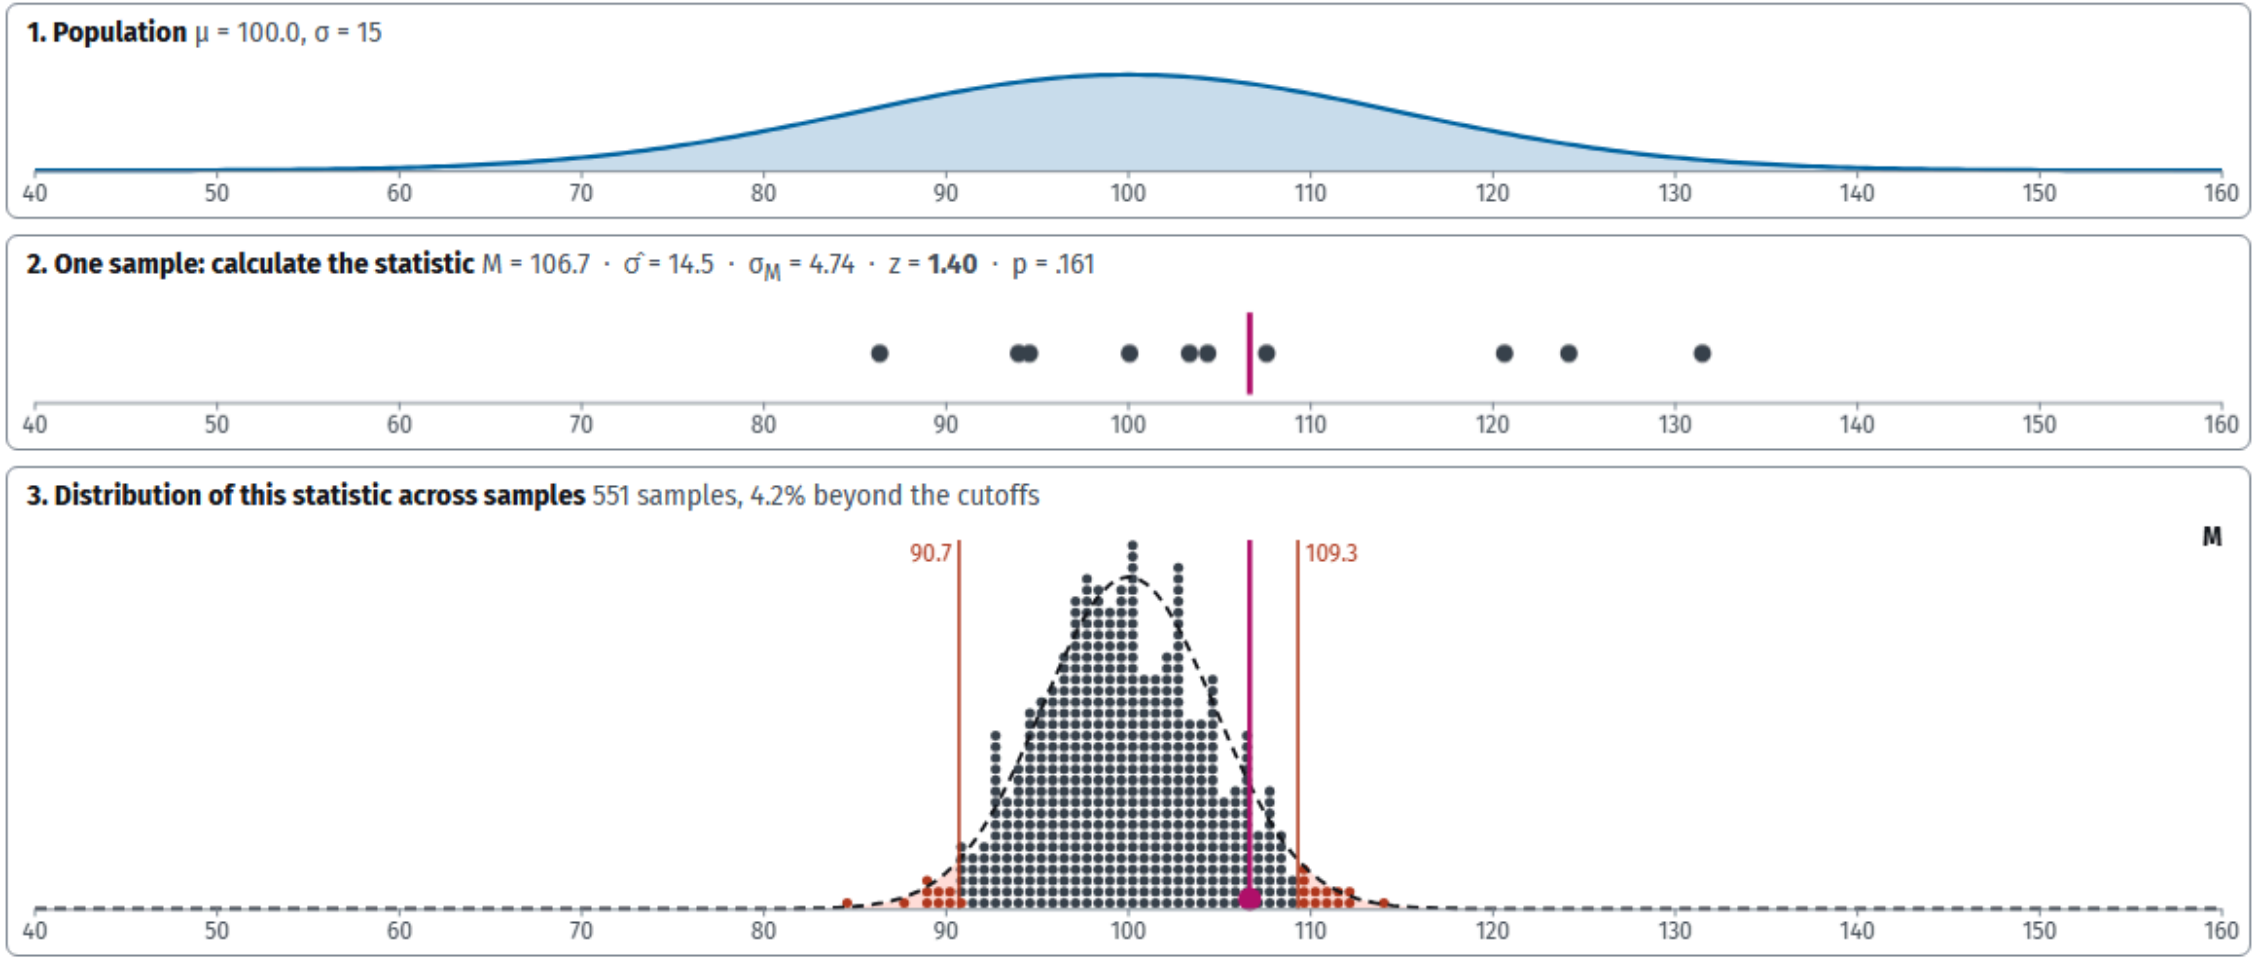
Simulator after Kristoffer Magnusson, rpsychologist.com.

Now each sample has five people



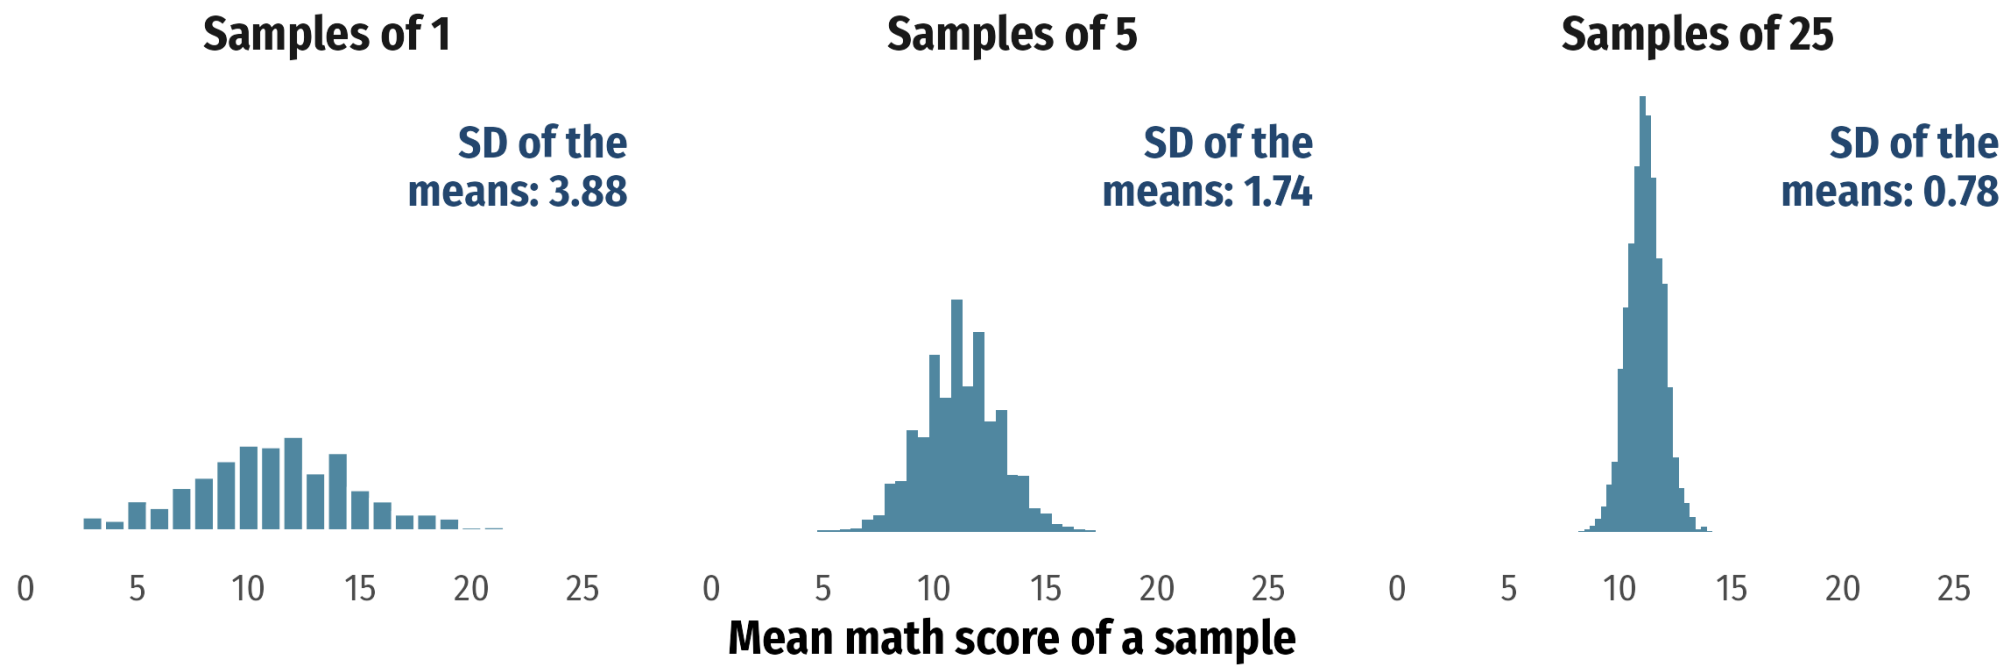
Simulator after Kristoffer Magnusson, rpsychologist.com.

With ten, the means get pointier again



Simulator after Kristoffer Magnusson, rpsychologist.com.

Our 619 students do the same thing



The law of large numbers: the bigger the sample, the closer M tends to be to μ

Data: Muralidharan, Singh, and Ganimian (2019). 5,000 samples of each size, treating the 619 students as a population.

That spread is the standard error

$$\sigma_M = \text{the standard deviation of all the sample means}$$

Sampling error

The gap for one sample: M − μ

Standard error

The typical gap, across all samples

With n = 1, σ_M is just σ. As n grows, σ_M shrinks

Predict it: How fast does it shrink?

Respond at **PollEv.com/adb**

Same population, with σ = 15. With samples of 1, σ_M = 15. Now each sample has 4 people. What's σ_M?

A 15

B 7.5

C 3.75

D 0

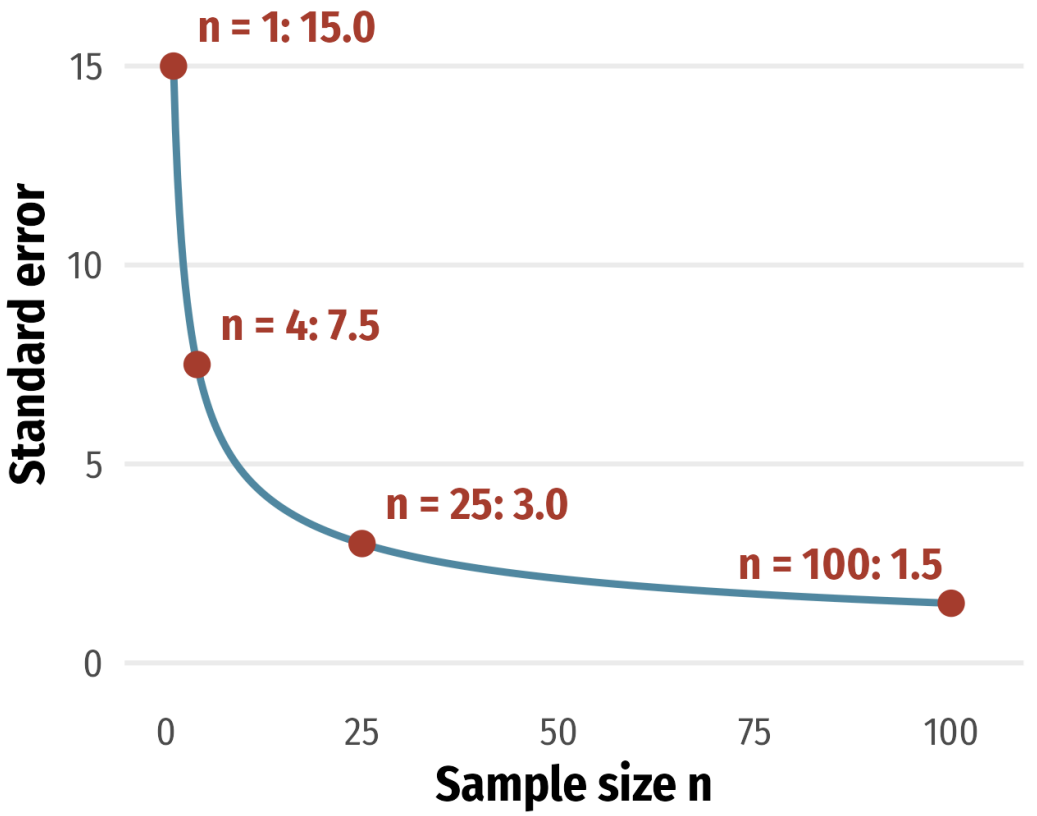
Estimating the standard error

Divide by the square root of n

$$\sigma_M = \frac{\sigma}{\sqrt{n}}$$

Our 619 students, n = 25: 3.90 / √25 = 0.78. The 5,000 samples gave 0.78

To halve the standard error, you need four times the people



Data: Muralidharan, Singh, and Ganimian (2019).

The whole central limit theorem

Shape

Close to normal as n grows, for any population

Center

The mean of all sample means is μ

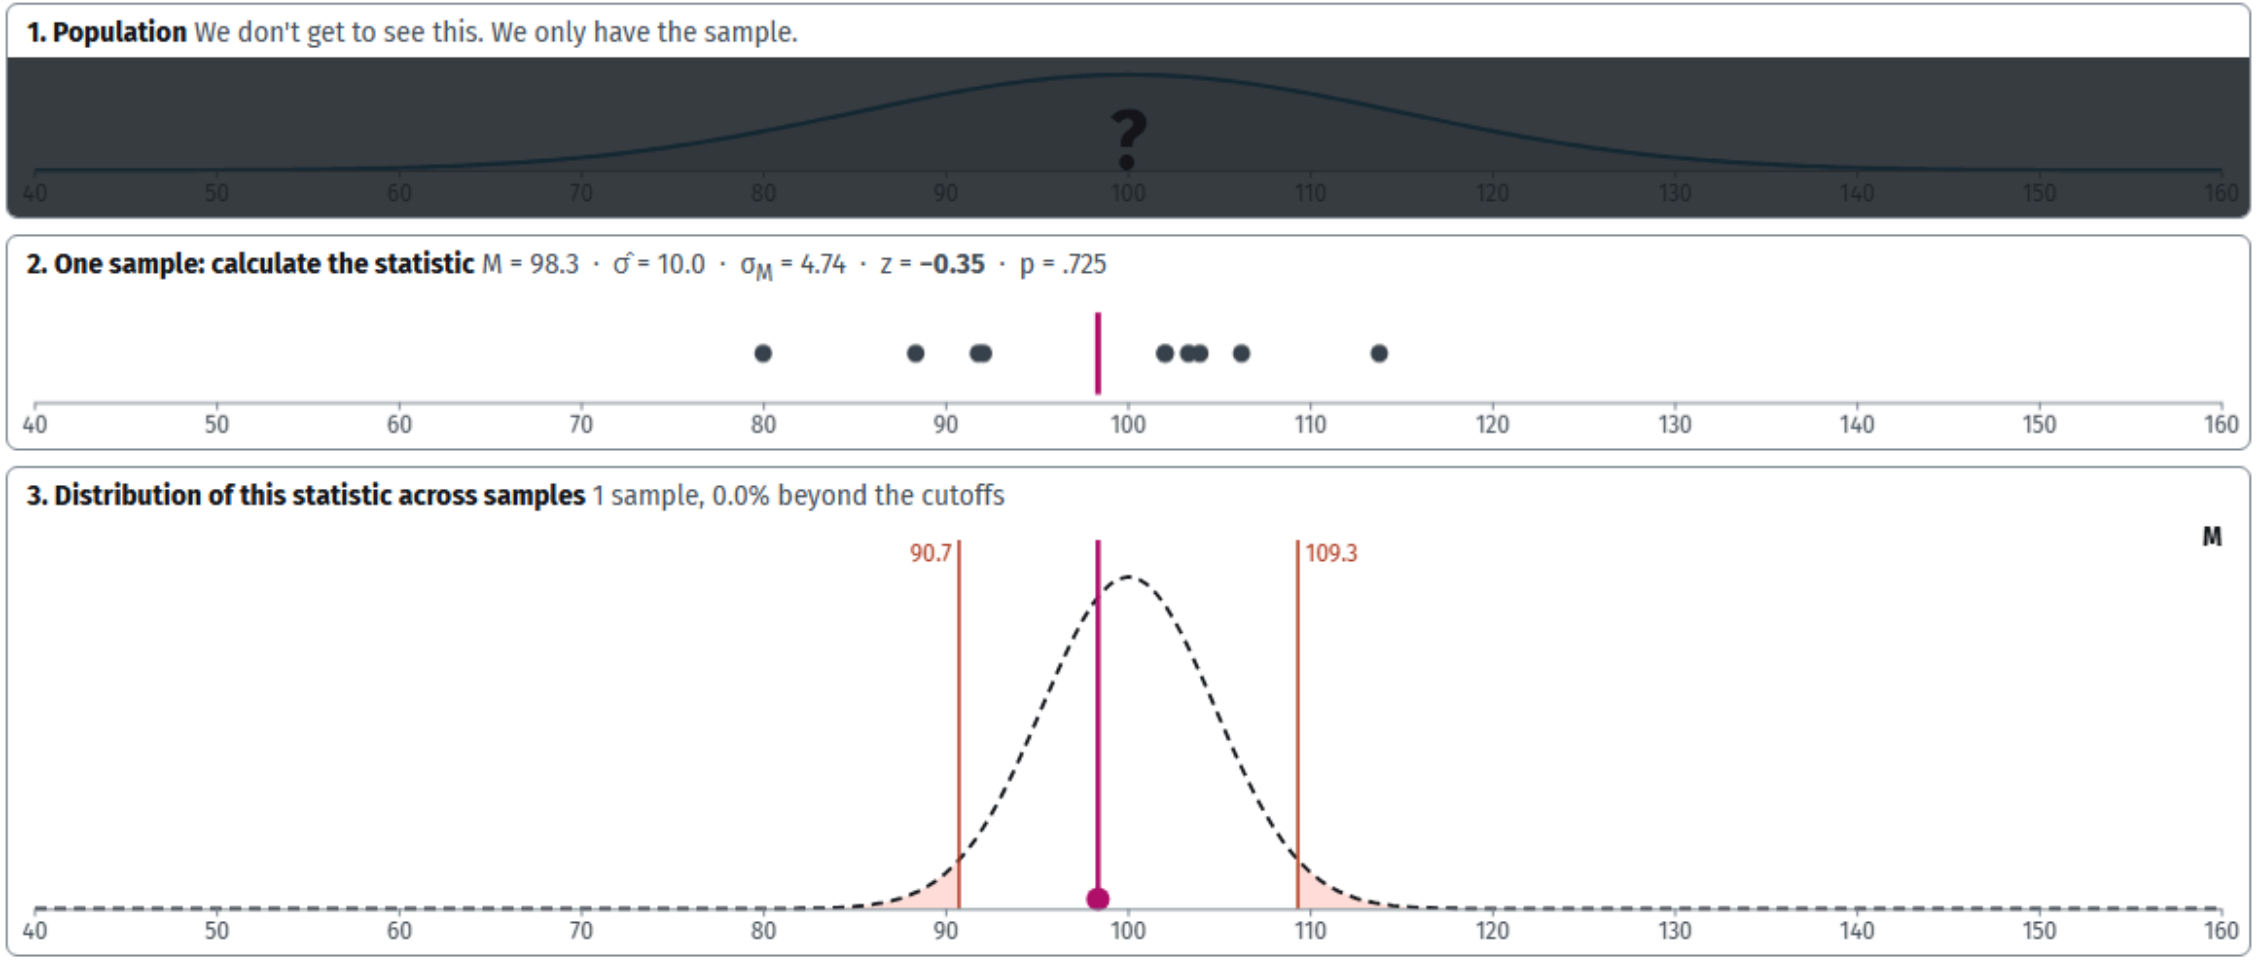
Spread

The SD of all sample means is $\sigma / \sqrt{n}$

Center and spread hold at any n. Only the shape needs a large n

In the book: the mean of all sample means is the *expected value of M*, and σ_M is the *standard error of M*

But we only ever get one sample



Simulator after Kristoffer Magnusson, rpsychologist.com.

So we estimate σ with our sample

$$\sigma_M = \frac{\sigma}{\sqrt{n}} \qquad \text{becomes} \qquad \hat{\sigma}_M = \frac{\hat{\sigma}}{\sqrt{n}}$$

Greek: the true standard error. It needs σ

Hat: our estimate of it, from one sample

In the book: $\hat{\sigma}$ is written s, and $\hat{\sigma}_M$ is written s_M. Chapter 9 introduces s_M

From scores to the standard error

Step	In symbols
1. Find each deviation from the mean	$X - M$
2. Square them and add them up	$SS = \sum(X - M)^2$
3. Divide by n - 1	$\hat{\sigma}^2 = SS/(n-1)$
4. Take the root	$\hat{\sigma} = \sqrt{\hat{\sigma}^2}$
5. Divide by the root of n	$\hat{\sigma}_M = \hat{\sigma}/\sqrt{n}$

Steps 1 to 4 are Topic 3. Step 5 is the only new one

The same thing, written other ways

$$\hat{\sigma}_M = \frac{\hat{\sigma}}{\sqrt{n}} = \frac{\sqrt{\hat{\sigma}^2}}{\sqrt{n}} = \sqrt{\frac{\hat{\sigma}^2}{n}} = \sqrt{\frac{\text{variance}}{n}}$$

The last form: divide the variance by n, then take the root

That form comes back in Topic 9, when we compare two means

One real sample of 25

	The truth (all 619)	Our one sample
Mean	$\mu = 11.12$	$M = 11.80$
Standard deviation	$\sigma = 3.90$	$\hat{\sigma} = 3.72$
Standard error	$\sigma_M = 0.78$	$\hat{\sigma}_M = 0.74$

One sample gives an estimate, and a sense of how far off it may be

Data: Muralidharan, Singh, and Ganimian (2019). One random sample of 25 from the 619 baseline scores.

Try it: Four real students

Math scores: 10, 8, 13, 13

Find M, SS, $\hat{\sigma}$ and $\hat{\sigma}_M$

Work alone for two minutes. Compare with a neighbor. Then we vote

Data: Muralidharan, Singh, and Ganimian (2019).

Try it: Answer

$M = (10 + 8 + 13 + 13) / 4 = 11$

$\hat{\sigma} = \sqrt{(18 / 3)} = \sqrt{6} = 2.45$

$SS = (-1)^2 + (-3)^2 + (2)^2 + (2)^2 = 18$

$\hat{\sigma}_M = 2.45 / \sqrt{4} = 1.22$

With four people, the mean could easily be off by a point or more

Data: Muralidharan, Singh, and Ganimian (2019).

Using standard errors

A standard error does two jobs

Job 1: Uncertainty

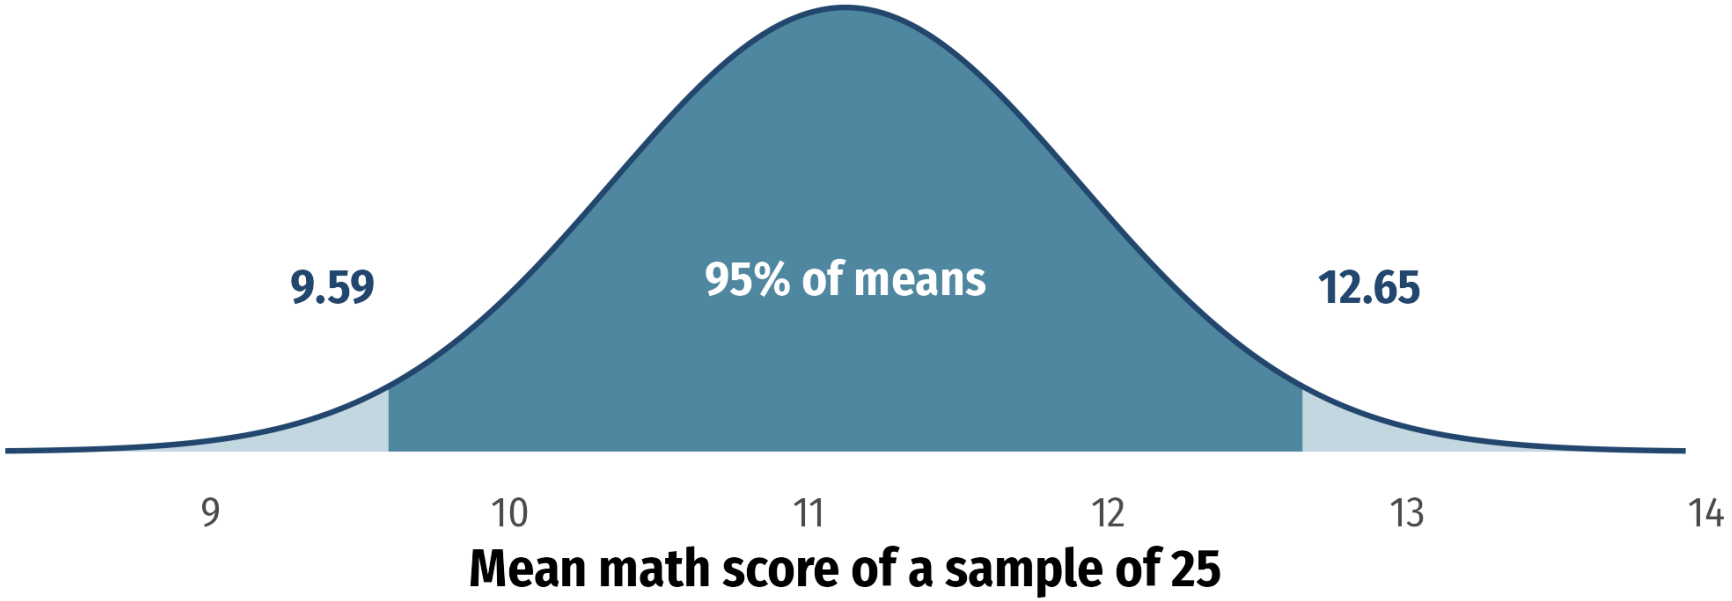
How far off might this estimate be?

Job 2: Surprise

If I'm right about the population, how weird is this mean?

The first job leads to confidence intervals. The second leads to hypothesis tests

Job 1: Where do means land?



$\mu \pm 1.96\ \sigma_M = 11.12 \pm 1.96 \times 0.78$. So 95% of samples land within 1.53 points of μ

Data: Muralidharan, Singh, and Ganimian (2019).

Polls call this the margin of error

A poll of 1,000 people: 52% say yes, ± 3 points

The ± is 1.96 standard errors

Here, 1.96 × 1.58 = 3.1 points

Sample size drives it, not population size

Four times the people: half the margin

Job 2: How surprising is this mean?

Locate a mean in its own sampling distribution, the way a z-score locates a score

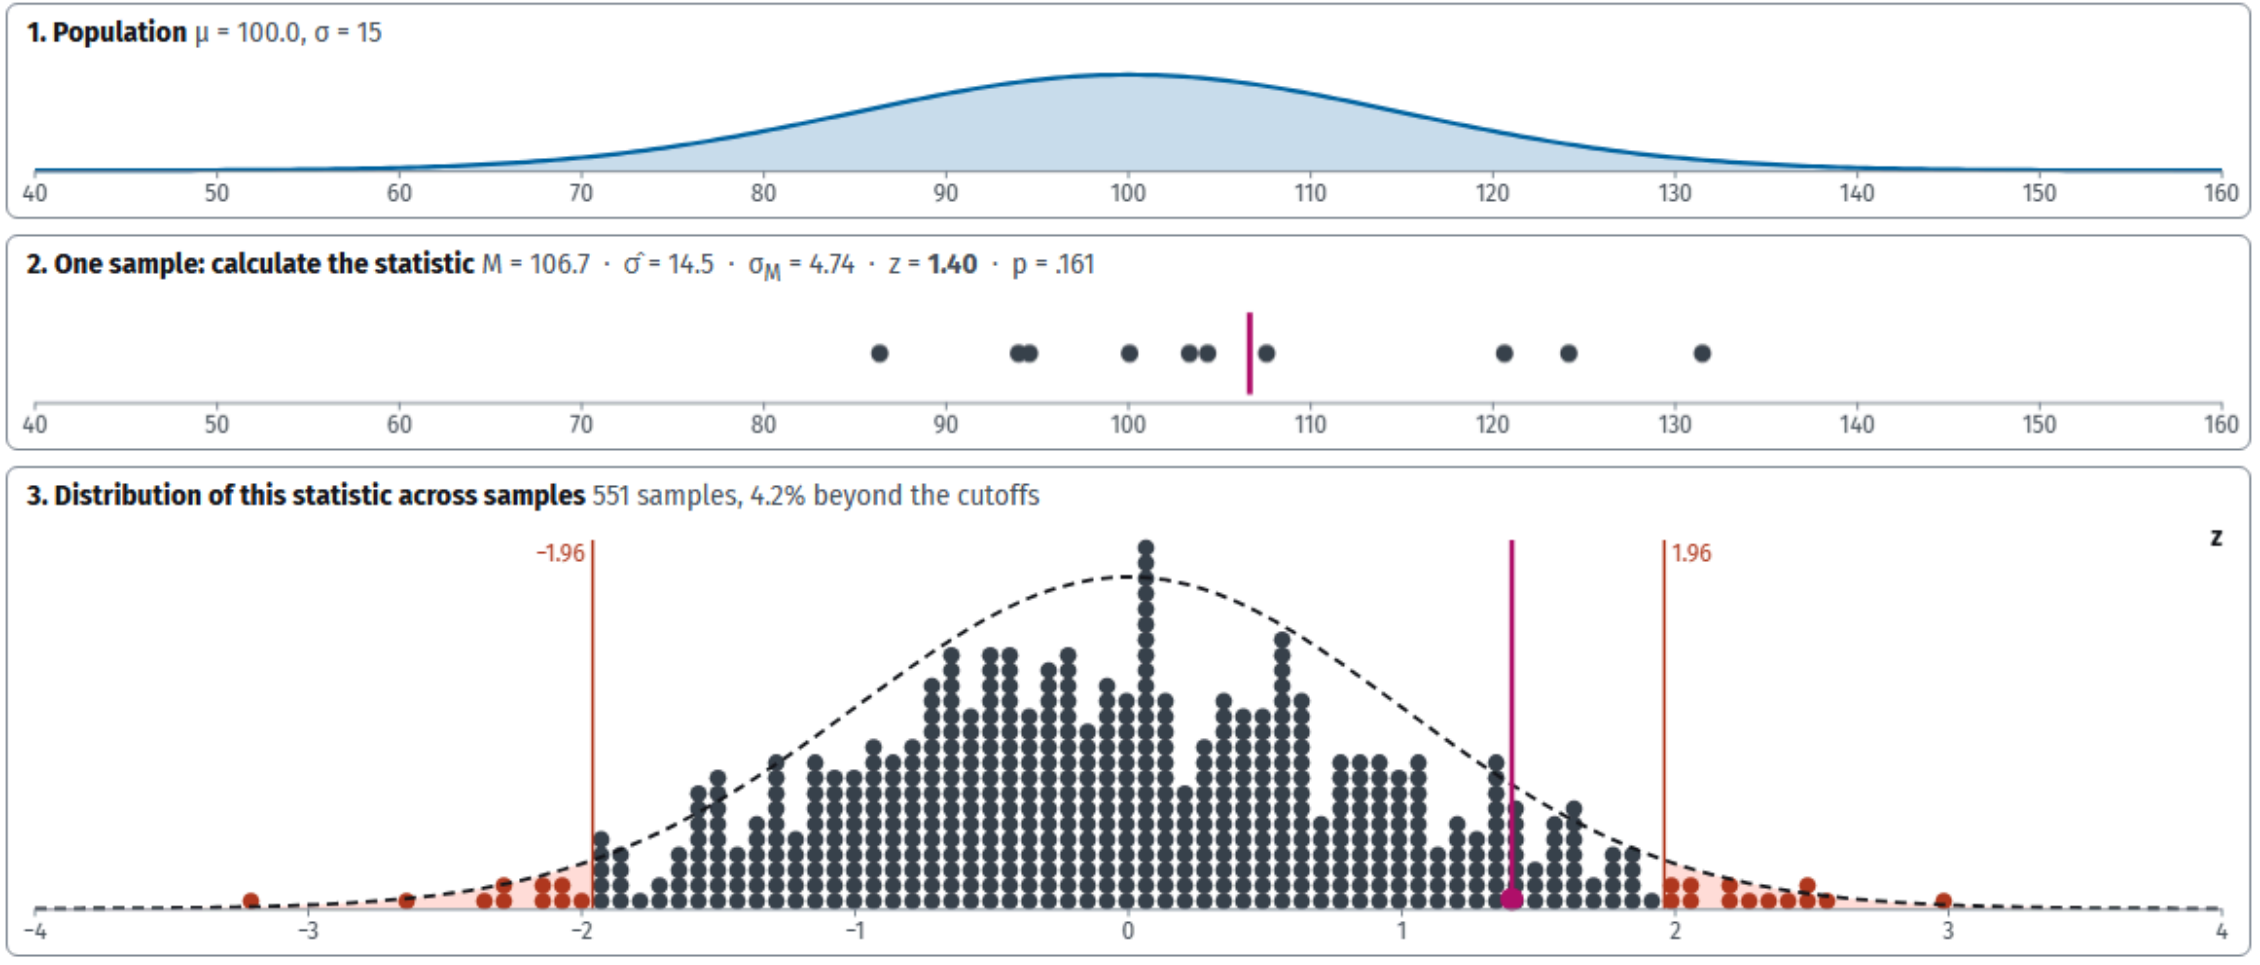
$$z = \frac{X - \mu}{\sigma} \qquad \text{becomes} \qquad z = \frac{M - \mu}{\sigma_M}$$

Top: how far M is from μ

Bottom: how far a mean typically falls

With n = 1, the two formulas are the same

In z-scores, every n lands on one curve



Simulator after Kristoffer Magnusson, rpsychologist.com.

Job 2 on real data: was the lottery fair?

The lottery picked 314 of the 619 students at random. So the winners are a random sample of our population

Baseline: $M = 10.984$, $\mu = 11.121$

$\sigma_M = 3.901 / \sqrt{314} = 0.220$

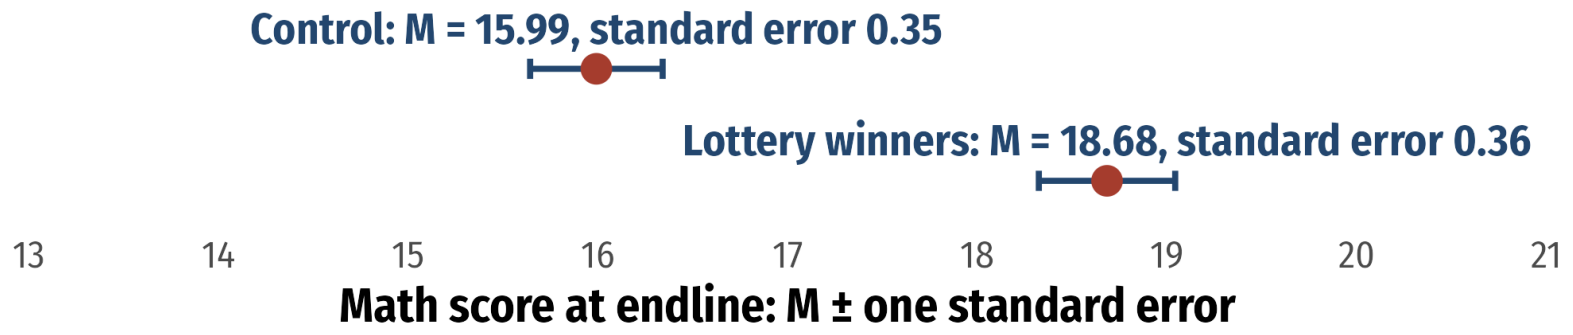
$z = (10.984 - 11.121) / 0.220 = -0.62$

53% of samples land this far from μ or farther

That's not weird at all, which is what a fair lottery should give

Data: Muralidharan, Singh, and Ganimian (2019).

Predict it: After the program



Could a gap like this be plain sampling error?

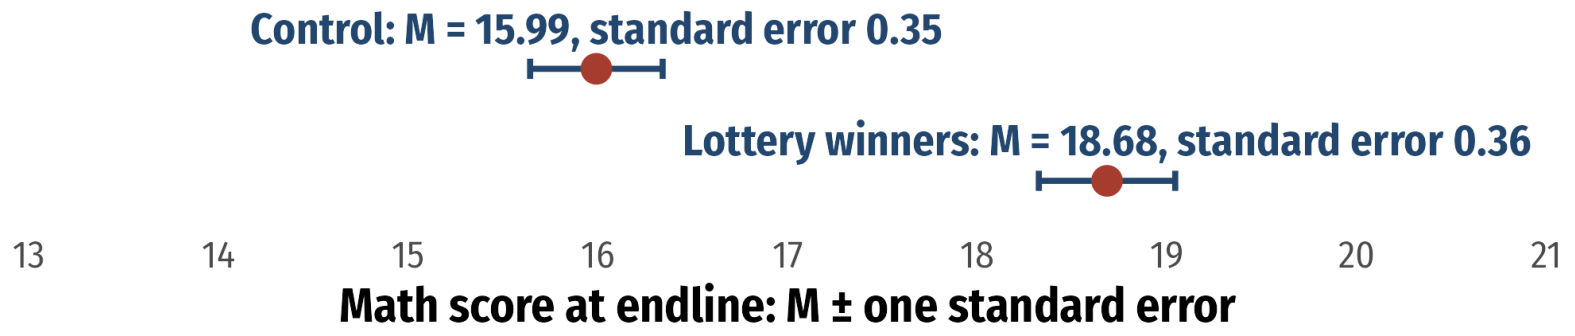
A Easily: both could come from one population

B Hardly: the gap is many standard errors

C No way to say without μ

Data: Muralidharan, Singh, and Ganimian (2019).

The gap dwarfs the sampling error



The gap is 2.69 points. Each mean is typically off by about 0.35 points

The answer is B. Topic 9 turns this into a real test, with a standard error for the gap itself

Data: Muralidharan, Singh, and Ganimian (2019).

Where are we?

Description

Summarize data

Visualize data

Prediction

We are moving here

Explanation

What you can do now

Read the blueprint: population, one sample, distribution of a statistic

Say what a sampling distribution is, and what each dot in it stands for

State the central limit theorem: shape, center and spread

Compute σ_M from σ, and $\hat{\sigma}_M$ from a sample

Use a standard error to find where means land, and a z-score for a mean

Next steps

Before Thursday

Read sections 8-1 to 8-3. Topic 7 is hypothesis testing

This Friday: Lab 4

Sampling distributions and the four steps. Bring a calculator and a laptop

Mastery and Problem Set for Chapter 7 are due Sunday, January 31 at 6 PM