Confidence Intervals

Keegan Korthauer

Department of Statistics

UW Madison

Point Estimation

- Estimate parameters with statistics (e.g sample mean estimates population mean)

- Single value (‘point’), never exactly equal to the true parameter

- How far off is it from the true parameter?

  - Uncertainty/Standard deviation gives us some idea

  - How confident are we that the true parameter lies in a certain interval?

Example - CD ROM Manufacturing

- Steel rods used in optical storage devices have a diameter specification of 0.45±0.02cm
- 1000 rods made in the last hour
- Need to indicate the percentage of acceptable rods in the population as an interval of the form **92%±x%**
- **x** is a number calculated to **provide reasonable certainty that the true population percentage is in the interval**.

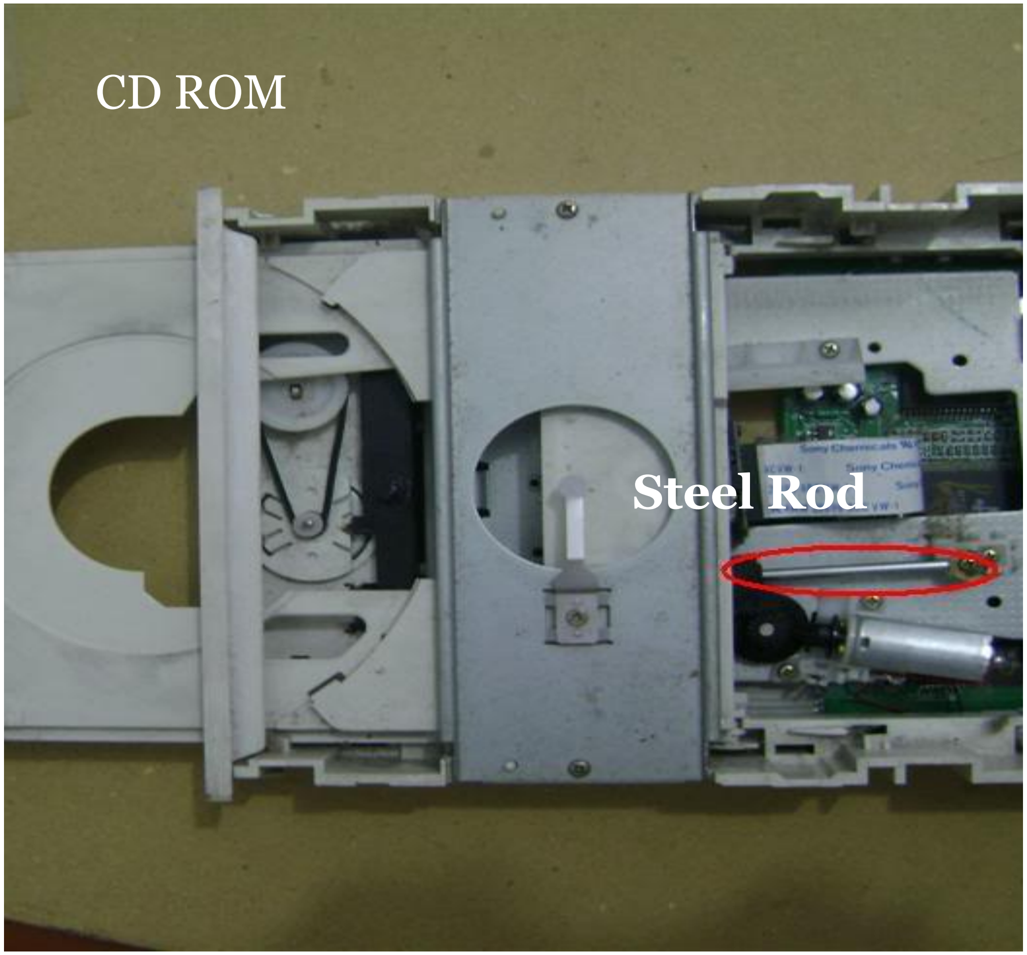

How to calculate x?

Another Example

- Consider the mean weight of 100 boxes $X_1, \ldots, X_{100}$ that have been filled with cereal by a certain machine on a certain day
- Sample mean $\overline{X}$ is 12.05 oz
- Sample standard deviation s is 0.1 oz
- Let the population mean and sd be μ and σ, respectively
- Since the sample is large, by the **CLT**, $\overline{X} \sim N(\mu, \sigma^2/100)$

Distribution of the Sample Mean - CLT

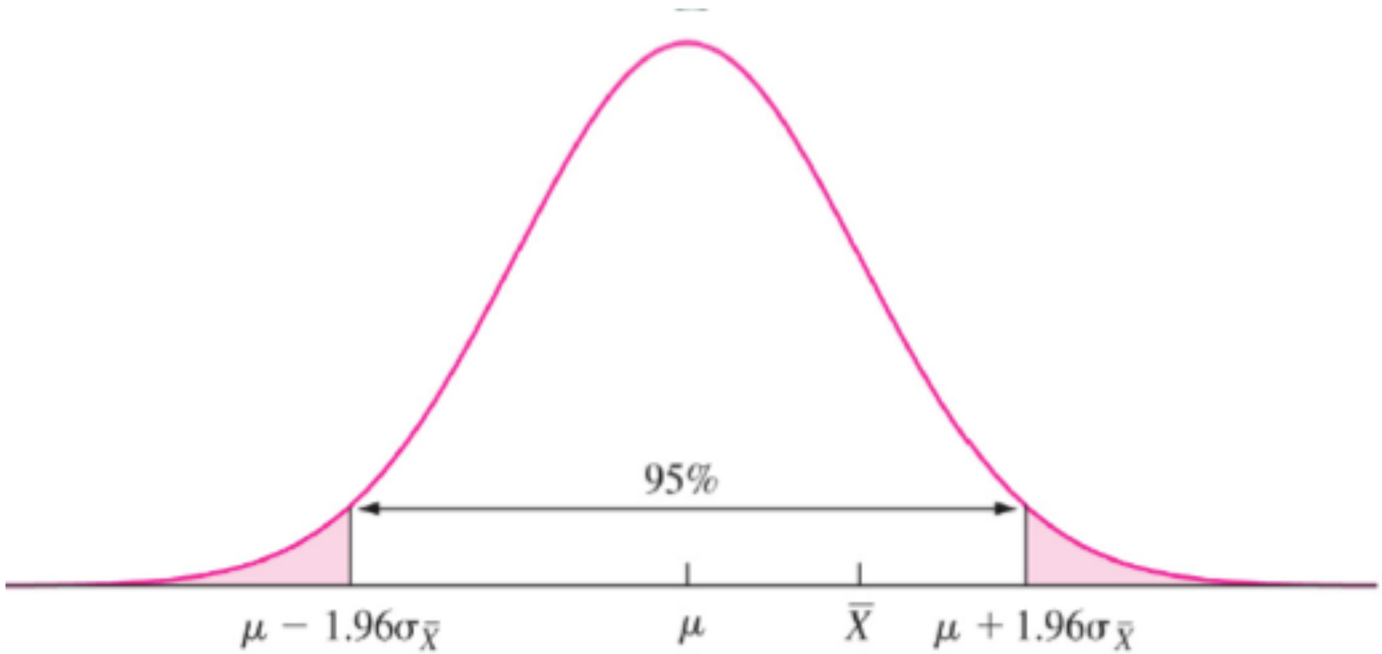

z	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
1.9	.9713	.9719	.9726	.9732	.9738	.9744	.9750	.9756	.9761	.9767

Where is the Observed Sample Mean?

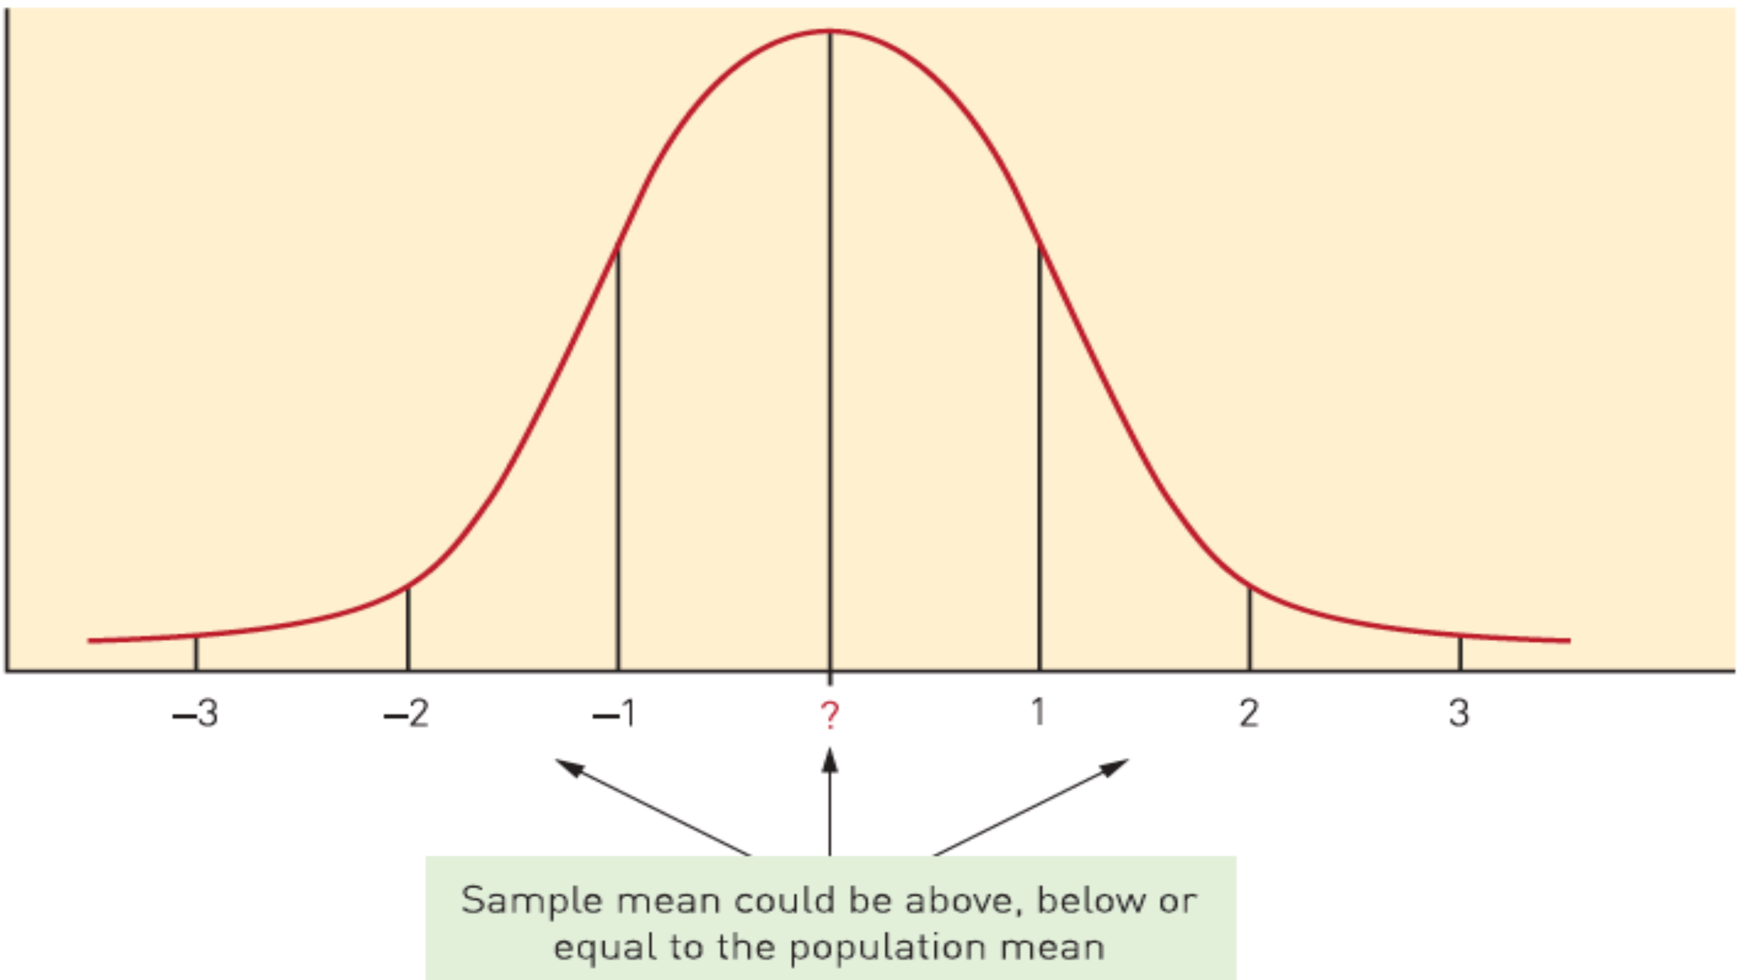

Population mean is **fixed**; sample mean is **random**

Confidence Interval for the Population Mean

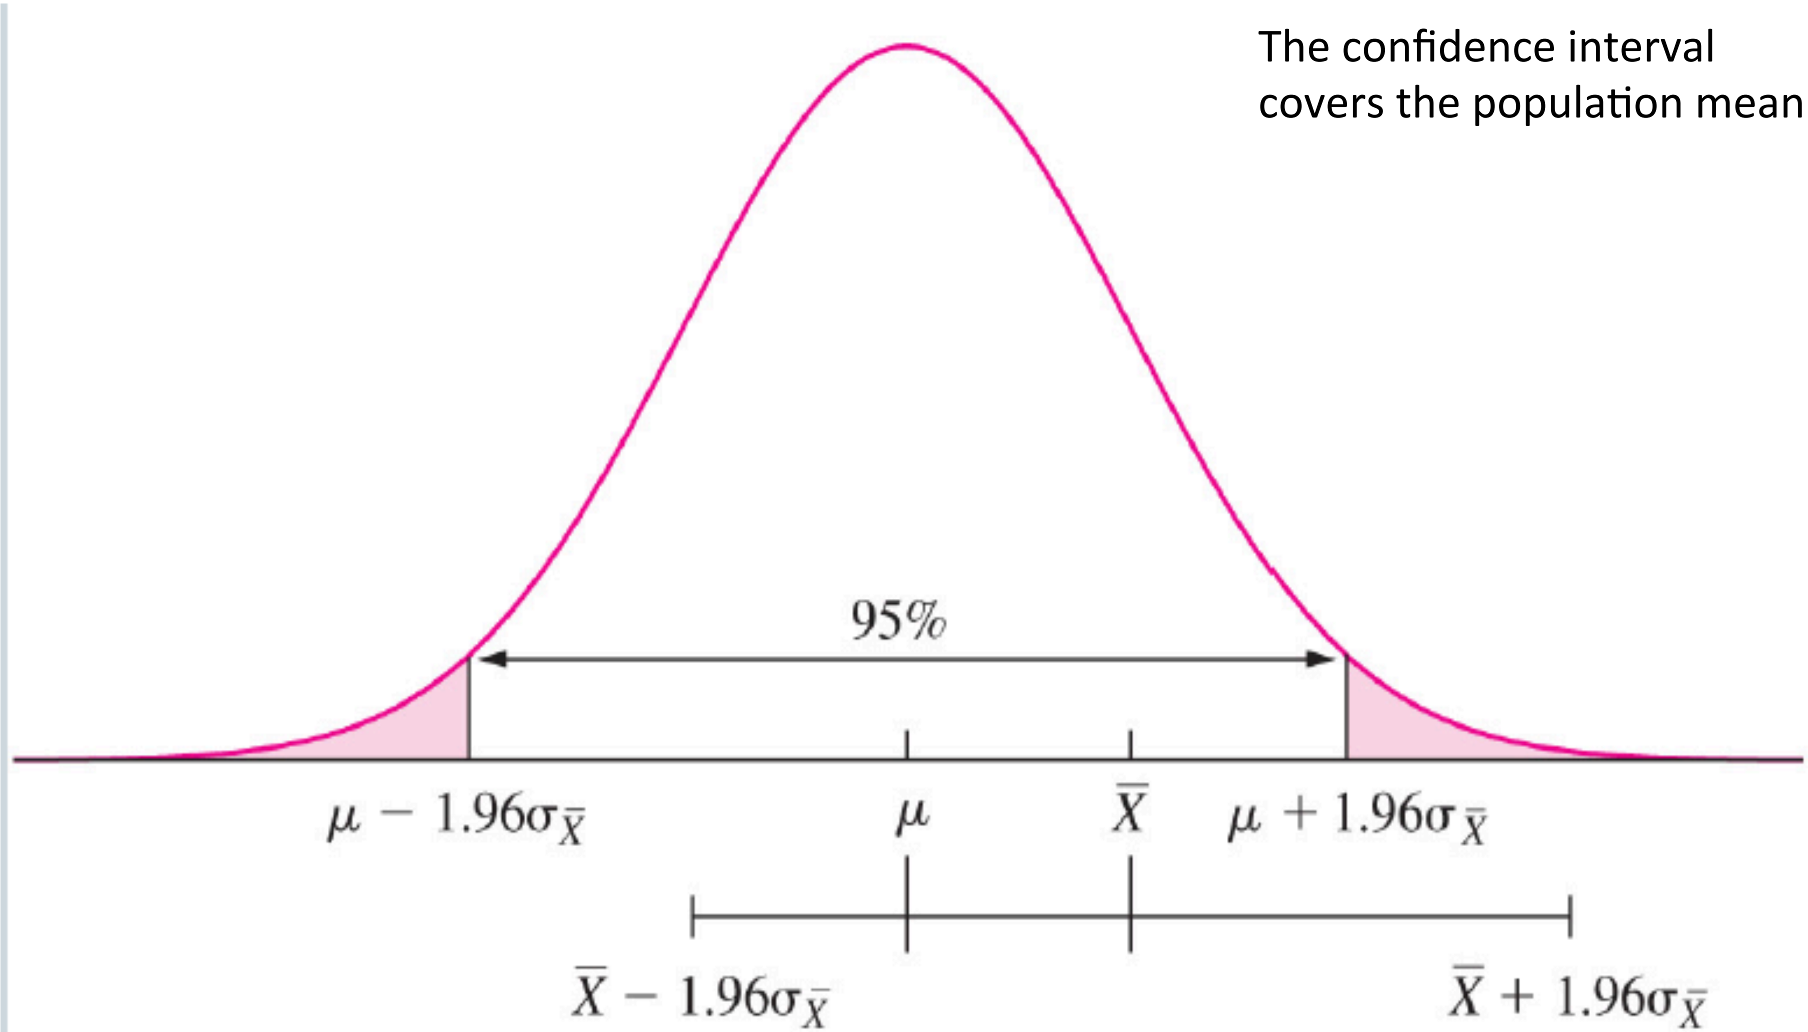

The confidence interval covers the population mean

95% confidence interval

Confidence Interval for the Population Mean

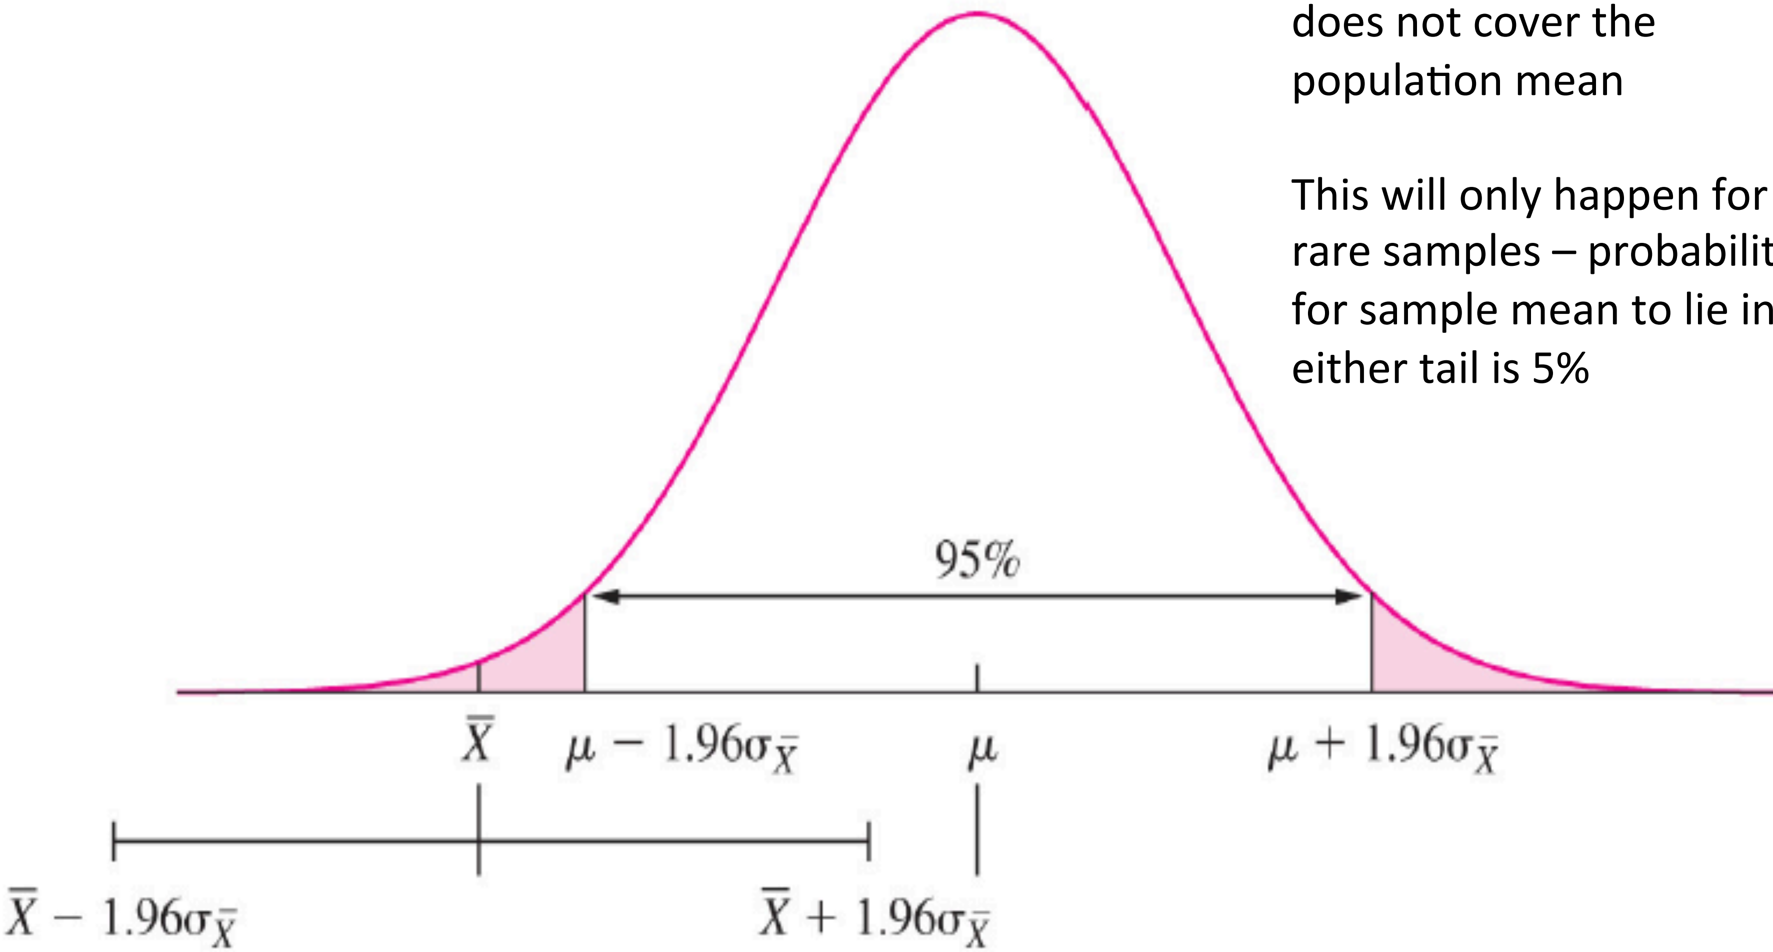

The confidence interval does not cover the population mean

This will only happen for rare samples – probability for sample mean to lie in either tail is 5%

95% confidence interval

Compute the Confidence Interval

In the cereal box example, we want to compute the 95% confidence interval (CI) for the population mean:

$$\bar{X} \pm 1.96\sigma_{\bar{X}}, \text{ where } \sigma_{\bar{X}} = \frac{\sigma}{\sqrt{100}}$$

We don't know σ but the sample is large so we can estimate it with s, which gives us

$$12.05 \pm 1.96 \frac{0.1}{\sqrt{100}} = (12.0304, 12.0696)$$

How to interpret?

Interpreting Confidence Intervals

Does this 95% CI actually cover the population mean μ?

- Depends on whether the sample mean came from the middle 95% of the distribution, or the outer 5%

- Since we don't know the true value of μ there is no way to tell for sure which portion of the distribution this particular sample came from

- If we construct these CIs on repeated samples, then 95% of the samples will have means in the middle 95% of the population and thus **95% of the confidence intervals will cover the population mean**

Interpreting Confidence Intervals

For this reason, for a 95% CI, we say that we have **95% confidence that the interval will cover the true population mean**

We use the term ‘confidence’ instead of probability because although the sample mean is random, the single interval we calculate is fixed

We also cannot talk about the probability that the population mean will lie within a certain interval, since it is also fixed

Confidence Levels

- We are not limited to 95% CIs (corresponds to 1.96 sds away from the mean)
- To obtain different levels of confidence need to use different Z-scores
  - Recall that for the normal distribution 68% of the population lies within 1 sd of the mean, so a 68% CI is:

$$\bar{X} \pm \frac{s}{\sqrt{n}}$$

  - Recall that for the normal distribution 99.7% of the population lies within 3 sds of the mean, so a 99.7% CI is:

$$\bar{X} \pm 3\frac{s}{\sqrt{n}}$$

Remember that μ is Fixed, sample mean is random!

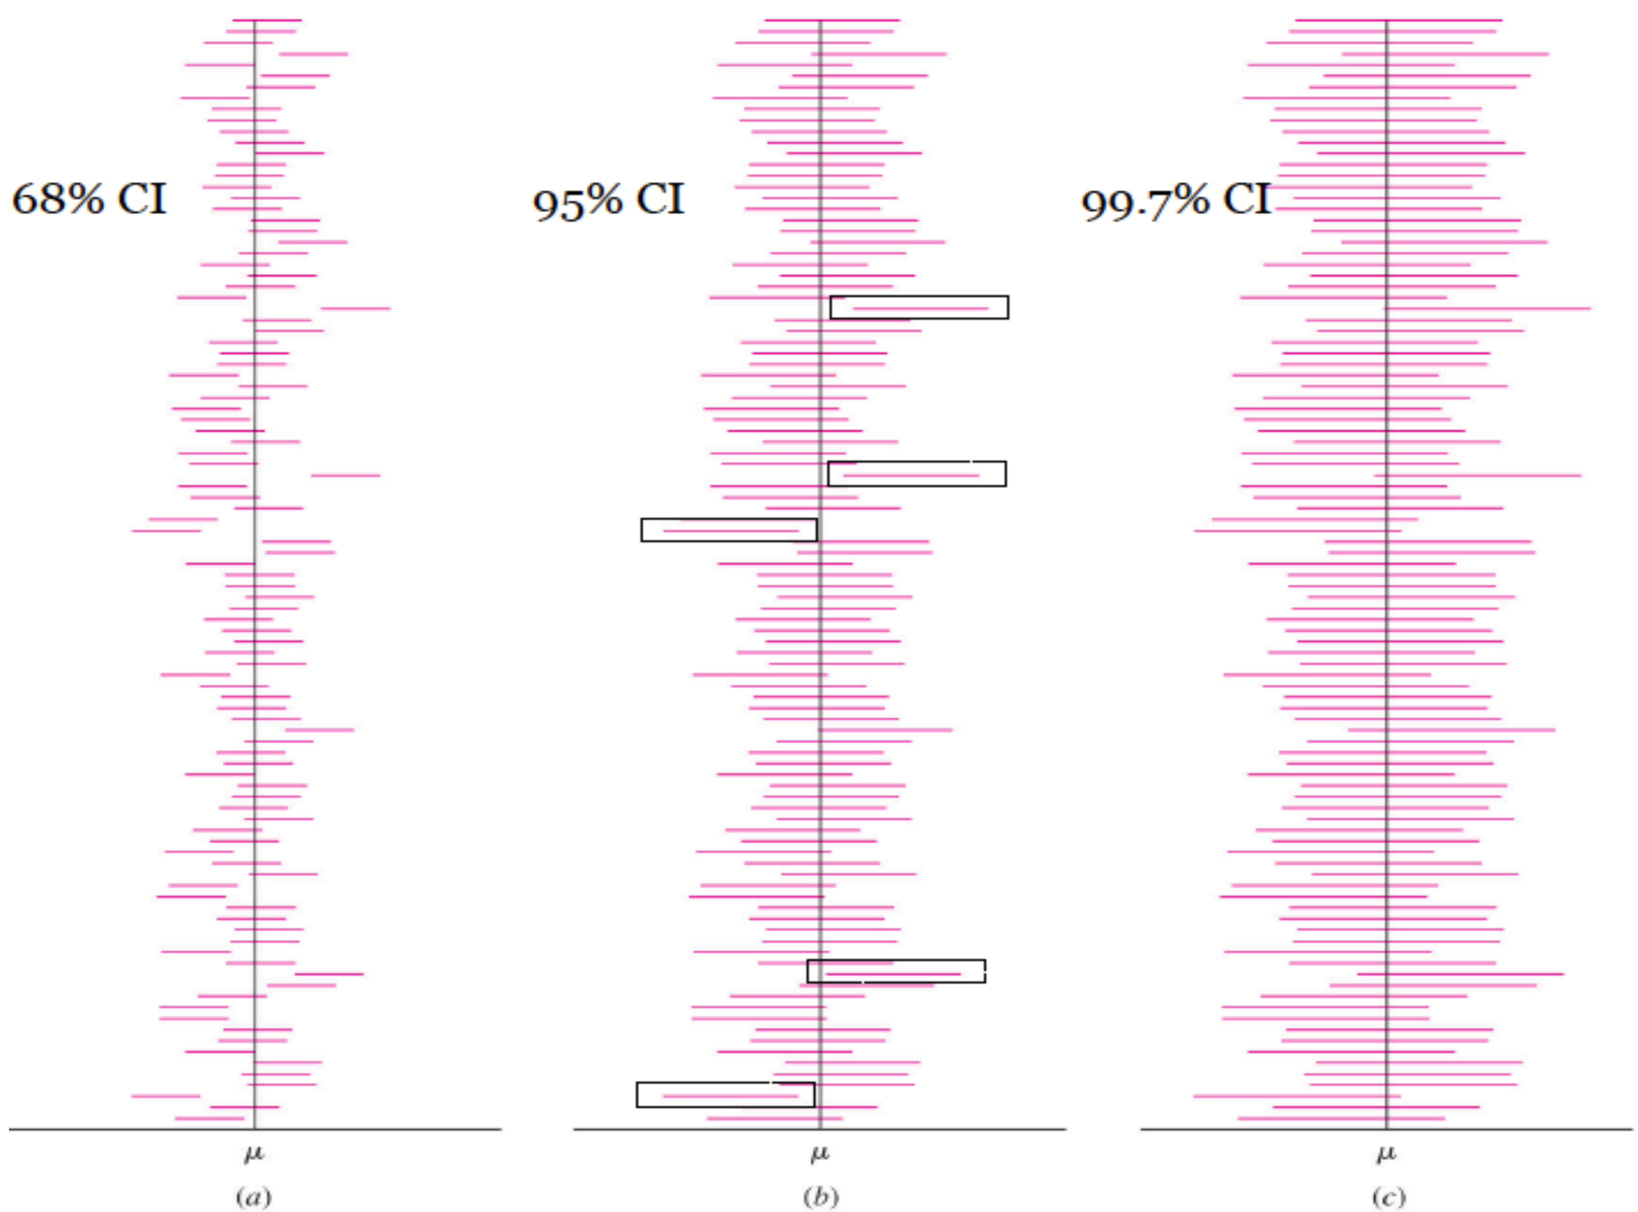

Confidence Intervals of Various Levels

- $\overline{X} \pm \frac{s}{\sqrt{n}}$ is a 68% confidence interval for μ.
- $\overline{X} \pm 1.645\frac{s}{\sqrt{n}}$ is a 90% confidence interval for μ.
- $\overline{X} \pm 1.96\frac{s}{\sqrt{n}}$ is a 95% confidence interval for μ.
- $\overline{X} \pm 2.58\frac{s}{\sqrt{n}}$ is a 99% confidence interval for μ.
- $\overline{X} \pm 3\frac{s}{\sqrt{n}}$ is a 99.7% confidence interval for μ.

Width Increases with Confidence

- The width of a CI will increase as we increase our level of confidence
- Two extreme cases:
  - we are 100% confident that the population mean will lie in the interval $(-\infty, \infty)$
  - we are 0% confident that the population mean will lie in the interval [12.05, 12.05]

Calculating a Confidence Interval of Arbitrary Level 100(1-α)%

Let $X_1,\ldots,X_n$ be a **large (n>30)** random sample from a population with mean μ and standard deviation σ, so that that $\bar{X}$ is approximately normal. Then a level **100(1-α)%** confidence interval for μ is:

$$\bar{X} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

When σ is unknown (usually the case), replace σ with the sample standard deviation s

For a 95% CI, α = 0.05.

How to find $z_{\alpha/2}$

$z_{\alpha/2}$ is the z-score that cuts off an area of α/2 in the right tail of the standard normal distribution

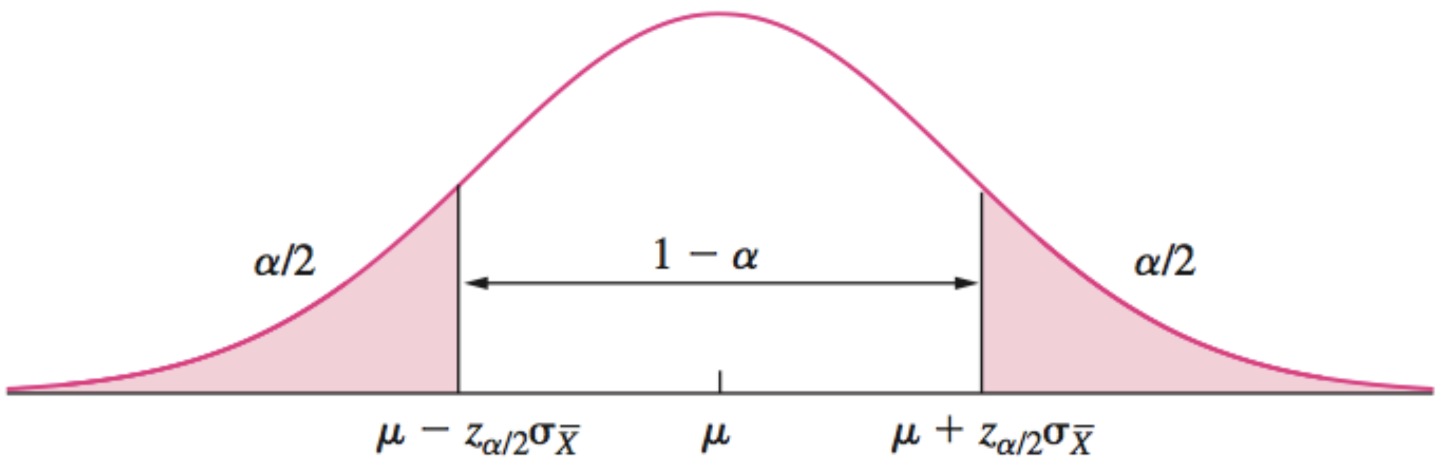

FIGURE 5.3 The sample mean $\overline{X}$ is drawn from a normal distribution with mean μ and standard deviation $\sigma_{\overline{X}} = \sigma/\sqrt{n}$. The quantity $z_{\alpha/2}$ is the z-score that cuts off an area of $\alpha/2$ in the right-hand tail. The quantity $-z_{\alpha/2}$ is the z-score that cuts off an area of $\alpha/2$ in the left-hand tail. The interval $\overline{X} \pm z_{\alpha/2}\sigma_{\overline{X}}$ will cover the population mean μ for a proportion $1 - \alpha$ of all samples that could possibly be drawn. Therefore $\overline{X} \pm z_{\alpha/2}\sigma_{\overline{X}}$ is a level $100(1 - \alpha)\%$ confidence interval for μ.

Example 5.1

The sample mean and standard deviation for the fill weights of 100 boxes are $\overline{X} = 12.05$ and $s = 0.1$. Find an 85% confidence interval for the mean fill weight of the boxes.

One-Sided Confidence Intervals

Also known as a confidence bound

- The level 100(1-α)% upper confidence bound for μ is

$$\bar{X} + z_{\alpha} \frac{\sigma}{\sqrt{n}}$$

- The level 100(1-α)% lower confidence bound for μ is

$$\bar{X} - z_{\alpha} \frac{\sigma}{\sqrt{n}}$$

Cereal Example Continued

The sample mean and standard deviation for the fill weights of 100 boxes are $\overline{X} = 12.05$ and $s = 0.1$.

Find a lower 95% confidence bound and a 99% upper confidence bound for the mean fill weight of the cereal boxes

Lower 95% confidence bound: 12.0336

Upper 99% confidence bound: 12.0733

Sample size for a given width

- Say we obtain a CI that is too wide to be useful
- As sample size increases, width decreases
- We want to know how many samples we need to obtain a CI of a certain width
- We need a sample of size n to obtain a CI of width *w* as given by the following equation:

$$n = \left(\frac{z_{\alpha/2} s}{w} \right)^2$$

rounding n up to the next largest integer

Example 5.8

In the fill weight example discussed earlier in this section, the sample standard deviation of weights from 100 boxes was $s = 0.1$ oz. How many boxes must be sampled to obtain a 99% confidence interval of width ± 0.012 oz?

Warning – Random Samples Only

The confidence interval method described so far only applies to random samples - bias or trends in the sample will violate the conditions of the CLT!

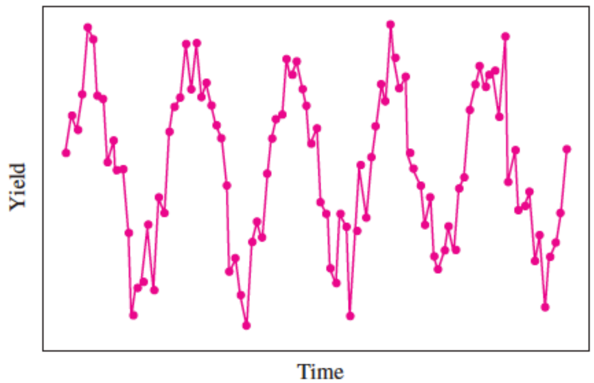

FIGURE 5.7 Yields from 100 runs of a chemical process, plotted against time. There is a clear pattern, indicating that the data do not constitute a random sample.

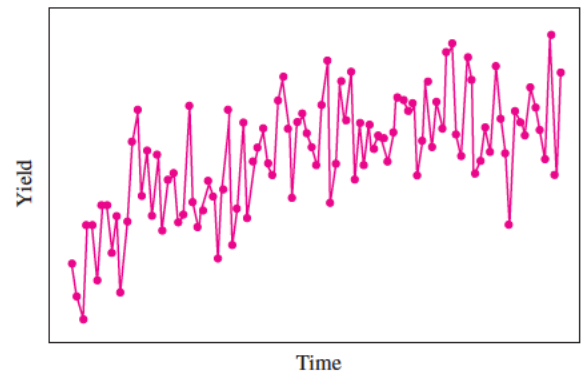

FIGURE 5.8 Yields from 100 runs of a chemical process, plotted against time. There is an increasing trend with time, at least in the initial part of the plot, indicating that the data do not constitute a random sample.

Next Time

- More on confidence intervals

- HW5 due on Friday