CHARTA LEARNING · LESSON NOTES

Samples, errors and confidence intervals

CFA Level I · Quantitative Methods

You will learn: Explain population and sample, tell sampling error from bias, state the central limit theorem, compute a standard error, build and read a confidence interval, and describe resampling.

Running example: Examples use made-up figures, such as 100 funds started ten years ago (70 survivors, 30 closed) and a fund with 36 monthly returns averaging 1.2% with a standard deviation of 3.0%.

1. Samples, sampling error and bias

- The population is every item you care about; the sample is the part you actually look at. A parameter describes the population (its true average); a statistic describes the sample. You see only the statistic, and use it to make a careful statement about the parameter.
- Sampling error is the sample statistic minus the true parameter. It is luck of the draw, not a blunder, and it shrinks as the sample grows.
- Probability sampling gives every member a known chance. Simple random: equal chances. Stratified random: split into groups (strata), then draw randomly from each in proportion to size. Cluster: pick whole natural clumps at random; cheap but less precise.
- Non-probability sampling relies on convenience (whatever is easy to reach) or judgement (an expert picks typical items). Sampling bias is a built-in tilt in how you pick; more data does not cure it. Survivorship bias is looking only at funds still alive.

100 FUNDS STARTED: 70 SURVIVORS AT 9.2% A YEAR, 30 CLOSED AT 3.2%		
Survivors	70 × 9.2	**644**
Closed funds	30 × 3.2	**96**
Average of all 100	(644 + 96) ÷ 100	**7.4%**
Survivorship bias	9.2 − 7.4	**1.8 points**
10,000 survivors instead?	still 9.2%	**bias stays**

Watch out: Treating bias like error: ask first how the sample was chosen, then how big it is.

2. Central limit theorem and standard error

- Take many samples of the same size and plot their averages: the sampling distribution of the mean. By the central limit theorem, averages of decent-sized samples form a bell shape (normal distribution), even if the population itself is lopsided. Rule of thumb: a sample size of about 30 or more.
- Three facts: the hill is centred on the true population mean, it is bell-shaped, and it narrows as the sample grows. It is the sample mean that becomes bell-shaped, not the data.
- The standard error is the typical distance between a sample mean and the true mean. If the population σ is unknown, use the sample standard deviation s.
- Four times the data halves the error; 100 times the data cuts it by 10; doubling cuts it only by about 30%.

Standard error = σ ÷ √n

STANDARD DEVIATION OF RETURNS 20%		
n = 25	20 ÷ √25 = 20 ÷ 5	**4%**
n = 100	20 ÷ √100 = 20 ÷ 10	**2%**
n = 400	20 ÷ √400 = 20 ÷ 20	**1%**

Watch out: Dividing by n instead of √n, and thinking the central limit theorem makes the data itself bell-shaped.

3. Confidence intervals

- A point estimate is one best-guess number; an interval estimate is a range around it. A confidence interval is a range likely to hold the truth, built from the guess, how sure you want to be, and the standard error.
- The sample mean $\bar{x}$ is the middle. The reliability factor z sets the confidence. z × standard error is the margin of error, stretched each side.
- Memorise z for two-sided intervals: 90% is 1.645, 95% is 1.96, 99% is 2.58. More confidence means a wider interval; more data means a narrower one.

Confidence interval = $\bar{x} \pm z \times (s \div \sqrt{n})$

N = 36, $\bar{X}$ = 1.2%, S = 3.0%, 95% (Z = 1.96)		
Standard error	3.0 ÷ √36 = 3.0 ÷ 6	**0.5%**
Margin of error	1.96 × 0.5	**0.98%**
95% interval	1.2 − 0.98 to 1.2 + 0.98	**0.22% to 2.18%**
99% (z = 2.58)	1.2 ± 1.29	**−0.09 to 2.49**
More data, n = 144, 95%	SE 3 ÷ 12 = 0.25; 1.2 ± 0.49	**0.71 to 1.69**

Watch out: Dividing by n, not √n, in the standard error.

4. Reading an interval, and z or t

- A 95% interval does not mean a 95% chance that the true mean is inside this one range. The true mean is fixed; the range is random because the sample is. The 95% describes the method: build 100 intervals and about 95 catch the truth.
- Use z when the spread is known or the sample is large; use t when the spread is unknown and the sample is small. The t distribution has fatter edges, so the interval is wider. As degrees of freedom grow, t approaches z.

Degrees of freedom = n − 1

N = 16, $\bar{X}$ = 8%, S = 12%, 95%, DF = 15		
Standard error	12 ÷ √16 = 12 ÷ 4	**3.0%**
t value (df = 15)	from t table	**2.131**
Margin of error	2.131 × 3	**6.39%**
95% interval	8 ± 6.39	**1.61 to 14.39**
If z = 1.96 (too tight)	8 ± 5.88	**2.12 to 13.88**

Watch out: Using z with a small sample and an unknown spread.

5. Resampling: bootstrap and jackknife

- Resampling treats your sample as a mini population, useful for statistics with no neat formula, such as a median. Bootstrap: draw n items with replacement (each goes back after a pick), compute the statistic, repeat many times. The spread of the answers estimates the standard error.
- Jackknife: leave out one observation at a time and recompute. A sample of n items gives exactly n answers, and always the same result.
- The bootstrap needs no formula or assumed shape and gives confidence intervals directly, but needs a lot of computing. It cannot invent information: a small or biased sample passes its flaws to every resample. It measures noise, not bias.

SAMPLE 2, 4, 6, 8, 10 (MEAN 6%), EIGHT RESAMPLES		
Resample 4, 4, 8, 10, 10	mean	**7.2**
Resample 2, 6, 6, 8, 4	mean	**5.2**
Resample 2, 2, 4, 6, 8	mean	**4.4**
Average of the 8 means	close to 6	**6.05**
Bootstrap standard error	spread of the 8 means	**1.41**
Formula check	3.16 ÷ √5	**1.41**

Key takeaways

- Sampling error is random and shrinks with more data; bias is systematic and does not.
- The central limit theorem: averages of big samples form a bell centred on the true mean.
- Standard error = s ÷ √n, so four times the data halves the error.
- Confidence interval = $\bar{x} \pm z \times SE$; higher confidence is wider, more data is narrower.
- The 95% belongs to the method, not to one interval. Use t for small samples with unknown spread.
- The bootstrap and jackknife measure error by resampling but cannot fix a biased sample.

Quick check

1. Sample of 100, standard deviation 10%. What is the standard error of the mean?
   Answer: 10 ÷ √100 = 1%.
2. What happens to an interval if confidence goes from 90% to 99%?
   Answer: It gets wider: z rises from 1.645 to 2.58.
3. A survivors-only fund database is made ten times bigger. Is the bias gone?
   Answer: No. Size cuts error, not bias.