SKILLSTACK SUITE

Statistics for Business & Decision Making — Lesson Workbook

All twelve modules in full: the idea, the method, the mistake to avoid, and the build task for each.

How to use this workbook

Twelve modules. Each one gives you the idea, the method, the mistake most people make, and a build task that produces a real artefact. Do the tasks in order — later modules assume the earlier artefacts exist. Reading without building will not get you the outcome.

Module 1 — What the question really is

What you will learn. Turning a business question into something measurable.

"Is marketing working?" cannot be answered. "Did trial signups per week increase after the March campaign, relative to the preceding eight weeks?" can. The translation step is where most analysis goes wrong — an unmeasurable question produces a number that feels like an answer but is not one.

The method

- Write the decision the answer will inform.
- Name the metric and the exact time window.
- Name the comparison — against what?
- Decide in advance what result would change the decision.

Most common mistake. Analysing before agreeing what would count as an answer.

Build task. Rewrite one vague business question as a measurable one.

Module 2 — Describing data

What you will learn. Centre, spread and shape — and when the average misleads.

The average is a summary of a distribution, and it is a bad one whenever the distribution is skewed. Revenue per customer, time-to-close and support ticket duration are almost always skewed, so the median usually describes the typical case better. Always look at spread and shape before quoting a centre.

The method

- Plot the distribution before quoting any summary.
- Report the median for skewed data, the mean for symmetric.
- Always quote a spread alongside a centre.
- Note the outliers rather than silently removing them.

Most common mistake. Quoting an average for a skewed distribution, where almost nobody sits near it.

Build task. Summarise a real dataset and state where the average misleads.

Module 3 — Visual honesty

What you will learn. Chart types and the ways graphs deceive.

Most misleading charts are not fraudulent, they are careless. A truncated y-axis exaggerates change; a dual axis manufactures correlation; a 3D pie makes comparison impossible. Being able to name the distortion protects you from believing your own charts as much as anyone else's.

The method

- Check the y-axis baseline first, every time.
- Be suspicious of any dual-axis chart.
- Match chart type to question type.
- Check whether the time window was chosen to flatter.

Most common mistake. Truncating the y-axis to make a modest change look dramatic.

Build task. Find a published chart and list its specific distortions.

Module 4 — Variation

What you will learn. Signal vs noise, and why last month's move is usually neither.

Every process varies. Reacting to normal variation as though it were signal produces constant, pointless change — the phenomenon behind most "the numbers moved, do something" meetings. Establish the normal range first, then only react outside it.

The method

- Plot the metric over many periods, not two.
- Establish the normal range of variation.
- React only to points outside that range.
- Look for sustained runs, not single moves.

Most common mistake. Comparing this month to last month and treating the difference as meaningful.

Build task. Plot one metric over 12+ periods and mark its normal range.

Module 5 — Sampling

What you will learn. What your data represents, and the bias in how it was collected.

Your data is a sample of something, even when it feels like everything. Survey respondents differ from non-respondents; logged users differ from churned ones. The question is never "is this biased" — it is "biased how, and does that bias affect this decision?"

The method

- Ask who or what is missing from this data.
- Ask how collection could systematically skew it.
- State the population you are generalising to.
- Decide whether the bias matters for this specific decision.

Most common mistake. Treating survey responses as representative of everyone you sent it to.

Build task. Write down the sampling bias in a dataset you already use.

Module 6 — Uncertainty

What you will learn. Confidence intervals and reporting a range.

A single number implies a precision you do not have. An interval communicates the uncertainty honestly and often changes the decision — a conversion rate of "4.2%" and "somewhere between 2.1% and 6.3%" support very different actions, even though the point estimate is identical.

The method

- Calculate the interval, not just the estimate.
- Report the range and the confidence level together.
- Note the assumptions the interval depends on.
- If the interval spans the decision threshold, you do not have an answer yet.

Most common mistake. Reporting a point estimate to two decimal places from a sample of thirty.

Build task. Put an interval around one number you currently report as exact.

Module 7 — Comparing groups

What you will learn. Hypothesis tests, p-values, and what significance does not mean.

A p-value is the probability of seeing data this extreme if there were no real difference. It is not the probability the hypothesis is true, and significance is not importance — with a large enough sample, a trivial difference becomes significant. Use the test to rule out chance, then use effect size to decide whether to act.

The method

- State the comparison before looking at the data.
- Choose the test that matches your data type.
- Interpret the p-value as evidence against chance, nothing more.
- Never run tests repeatedly until one comes out significant.

Most common mistake. Saying "p < 0.05 so the hypothesis is 95% likely to be true". It says nothing of the kind.

Build task. Run a comparison test on two real groups and interpret it in plain language.

Module 8 — Effect size

What you will learn. Whether a real difference is worth acting on.

Significance tells you a difference is probably real. Effect size tells you whether it matters. A 0.3% lift that is statistically bulletproof may not cover the cost of the change. Always pair the two, and always convert the effect into business units before recommending anything.

The method

- Calculate the effect size, not just the p-value.
- Convert it into money, time or customers.
- Compare that against the cost of acting.
- Report both together, always.

Most common mistake. Recommending action on significance alone, without sizing the effect.

Build task. Calculate the effect size and state whether it justifies acting.

Module 9 — Relationships

What you will learn. Correlation, regression and confounders.

Correlation measures how two things move together; regression estimates how much one moves with the other. Neither establishes cause. The confounder — a third variable driving both — is the usual explanation for a surprisingly strong business correlation.

The method

- Plot the relationship before calculating anything.
- Check for a third variable that could drive both.
- Treat regression coefficients as associations, not effects.
- Watch for a relationship that only exists in one subgroup.

Most common mistake. Reading a regression coefficient as "if we increase X, Y will rise by this much".

Build task. Find a correlation in your data and name a plausible confounder.

Module 10 — Causation

What you will learn. What it takes to claim cause.

Cause requires either a controlled experiment or a natural experiment with a credible comparison group. Observational business data almost never supports a causal claim, however large the dataset. The honest phrasing is "associated with", and using it protects you when the decision is questioned.

The method

- Ask whether anything was randomised.
- If not, ask what the comparison group is and why it is credible.
- Check the timing — did the cause precede the effect?
- Default to "associated with" unless you ran an experiment.

Most common mistake. Claiming cause from a before-and-after comparison with no control.

Build task. Rewrite one causal claim in your reporting as an associational one.

Module 11 — Experiments

What you will learn. Designing an A/B test that answers the question you asked.

A valid test fixes its parameters before it starts: the metric, the sample size, the duration and the threshold for acting. Deciding any of these after seeing data turns the test into a search for a favourable stopping point.

The method

- Fix the primary metric before launching.
- Calculate the sample size needed for the effect you care about.
- Fix the duration and do not stop early on a good result.
- Write the decision rule before you look at anything.

Most common mistake. Peeking at results daily and stopping when it looks good. This manufactures false positives.

Build task. Design an A/B test for a decision your team actually faces.

Module 12 — Communicating results

What you will learn. Presenting findings with the uncertainty intact.

The temptation is to strip caveats so the recommendation lands harder. That works until the decision is questioned and the caveats surface anyway. Lead with the answer, attach the range, name the main assumption, and state what would change your mind.

The method

- Lead with the decision-relevant answer.
- Attach the uncertainty range to the headline number.
- Name the single biggest assumption.
- State what evidence would change the conclusion.

Most common mistake. Stripping the uncertainty to sound more confident.

Build task. Present one finding with the range and assumptions intact.