THE GRIPHAN ONLINE SCHOOL

UNIVERSITY LIBRARY

Statistics for Researchers II — Data Analysis

From samples to conclusions: estimation, testing and relationships

Nji Griphan

griphan.48hourrfprescue.com

Statistics for Researchers II — Data Analysis

From samples to conclusions: estimation, testing and relationships

Nji Griphan

The Griphan Online School — University Library

IMPRINT

This book is published for the Library of the Griphan Online School. Author and publisher of record: **Nji Griphan**. School address on the web: **griphan.48hourrfprescue.com**. Professional mailbox: **info@griphan.48hourrfprescue.com**. © 2026 Nji Griphan. This book was written for the students of this school. Students enrolled at the school may download and keep it for their own study. The library grows shelf by shelf until it carries the whole university.

A note on this book

Volume I taught you to describe the data in hand. This volume teaches the harder step: reasoning from a sample to the population it came from — estimation with confidence intervals, significance tests chosen and reported honestly, and the analysis of relationships between variables, ending with the design of studies that

can support causal claims.

Contents

Chapter 1 — Sampling and the Sampling Distribution

Chapter 2 — Estimation and Confidence Intervals

Chapter 3 — Significance Tests — Logic and Limits

Chapter 4 — Comparing Groups — t-tests and Chi-square

Chapter 5 — Relationships — Correlation and Regression

Chapter 6 — Designing Studies That Can Conclude

Appendix — Worked Examples and Practice

Glossary of Key Terms

Chapter 1

Sampling and the Sampling Distribution

Inference works because samples, though variable, vary in a knowable pattern. This chapter explains sampling variation, the standard error, and why the same study repeated would not give the same number twice — and why that is manageable rather than fatal.

Samples vary

Draw two random samples of 50 households from one village and their mean incomes will differ — a little, usually. Sampling variation is not error or failure; it is the predictable consequence of observing a part instead of the whole, and statistics exists to measure it rather than pretend it away.

The standard error

The standard error is the standard deviation of a sample statistic across the samples you might have drawn. For a mean it falls as sample size rises — with the square root of n, so quadrupling the sample only halves the error. This square-root law is why precision is expensive and why enormous samples are rarely the best use of research money.

The central limit pattern

Averages of samples cluster in a bell-shaped pattern around the true population mean, whatever the shape of the underlying data, once samples are moderately large.

This central limit pattern is the engine under confidence intervals and most common tests; it is a fact about averages, earned by mathematics and confirmed by simulation.

Design still rules

No formula repairs a biased sample. Inference assumes the sample was drawn by a method that gave the population a fair chance to appear. The field discipline of Volume I — defined populations, honest sampling — is the foundation on which everything in this volume stands.

Key points

- Sampling variation is predictable, not shameful.
- Standard error shrinks with the square root of sample size.
- Inference cannot repair a biased sample.

Exercises

1. A study's standard error is 8 with n = 100. What n halves it?

2. Explain the central limit pattern in plain words to a farmer cooperative.

3. Why does doubling a biased sample not improve its conclusions?

Chapter 2

Estimation and Confidence Intervals

A sample gives not the population value but a neighbourhood where it plausibly lives. This chapter teaches you to build and read confidence intervals — the honest alternative to pretending a sample number is exact.

Point and interval estimates

The point estimate is the sample value itself: mean yield 1,840 kg/ha. The interval estimate adds the margin: 1,840 ± 160, the 95 per cent confidence interval. Reporting the interval is reporting both your finding and your uncertainty in one breath.

What 95 per cent means

A 95 per cent confidence interval is built by a method that captures the true value in 95 of every 100 such studies. It does not mean there is a 95 per cent chance this particular interval holds the truth — the truth is fixed; it is the method that succeeds 95 times in 100. The distinction sounds fine and matters in honest reporting.

Width is information

A wide interval confesses a small sample or scattered data; a narrow one claims precision. Never apologise for width — it is the data speaking. Do distrust intervals that are suspiciously narrow for the fieldwork described; arithmetic, like axes, can be wounded.

Intervals for proportions

Proportions carry intervals too: 46 of 200 households food insecure gives 23 per cent with a margin near ±6 points. Beside every headline percentage in your reports, the interval must stand — it is the difference between a finding and a slogan.

Key points

- Report the interval, not only the point.
- 95 per cent describes the method's long-run success.
- Interval width honestly reports your uncertainty.

Exercises

1. A mean is 52 with margin ±5 at 95%. Write the interval and its plain meaning.

2. Why is '23% ± 6' more informative than '23%'?

3. A proportion uses n = 25. Predict, and explain, the interval's width.

Chapter 3

Significance Tests — Logic and Limits

A significance test asks a narrow question: if there were really no effect, how surprising would data like ours be? This chapter teaches that logic, the p-value's true meaning, and the abuses that fill journals and newspapers alike.

The null hypothesis

Every test sets up a null hypothesis — no difference, no effect — and asks whether the data are strange enough, under that assumption, to doubt it. The test never proves the alternative; it measures the discomfort of the null. Keeping this straight prevents nine of ten misreadings.

The p-value, honestly

The p-value is the probability, computed under the null, of results at least as extreme as those observed. p = 0.03 does not mean a 3 per cent chance the null is true, nor a 97 per cent chance your effect is real. It is a measure of surprise, nothing more — useful, and routinely oversold.

Errors of both kinds

A test can cry wolf — declaring an effect that is not there (Type I) — or miss a real effect (Type II), especially in small samples. The 0.05 convention merely fixes the wolf rate at one in twenty. With twenty tests, expect one false alarm by design; researchers who test until something shines are manufacturing wolves.

Significance is not importance

In a huge sample, a trivial difference becomes 'significant'; in a small one, a vital difference may not. Always report the effect's size with its interval beside the p-value. The reader deserves to judge importance; the test only reports surprise.

Key points

- Tests measure surprise under the null — they never prove the alternative.
- p is not the probability that the null is true.
- Report effect size and interval with every p-value.

Exercises

1. Rewrite 'p = 0.04 proves the treatment works' honestly.

2. A researcher runs 20 tests and reports the one with p = 0.03. Critique.

3. Give an example of a significant-but-trivial and an important-but-nonsignificant result.

Chapter 4

Comparing Groups — t-tests and Chi-square

Two workhorses cover most comparisons a researcher will meet: the t-test for means and the chi-square test for counts in categories. This chapter states what each assumes, when each applies, and how to report both.

Comparing two means

The t-test compares two group means relative to the scatter within groups and the group sizes. Independent samples — different plots, different households — use the standard form; paired data — the same farmers before and after — use the paired form, which is stronger because each unit serves as its own control.

Assumptions, plainly

The t-test assumes independent observations, roughly bell-shaped data (or large samples), and honest measurement. It tolerates modest departures; it does not tolerate dependent data treated as independent — weighing the same ten plants twice does not make twenty observations.

Counts in categories — chi-square

When outcomes are categories — adopted or not, passed or failed — arrange counts in a table and ask chi-square whether the distribution differs between groups more than chance explains. Report the table itself; a test on counts the reader cannot see is a claim on trust.

Choosing the test

Means of numerical outcomes: t-test family. Counts in categories: chi-square. Paired versus independent decides the variant. When in doubt, state the data's shape first — the right test is usually the one the data's shape points to.

Key points

- t-test for means; chi-square for counts in categories.
- Paired designs are stronger: each unit is its own control.
- Always print the table or the group summaries beside the test.

Exercises

1. Choose tests: (a) yields of two varieties on separate plots; (b) pass/fail by study group; (c) weights of the same calves at two ages.

2. Why do ten plants weighed twice not give n = 20?

3. What must accompany a chi-square result in a report?

Chapter 5

Relationships — Correlation and Regression

Research questions often ask how variables move together — rainfall with yield, study hours with scores. This chapter teaches correlation as a measured association and regression as its working tool, with the causal warning posted at every door.

Correlation

The correlation coefficient r runs from −1 to +1: sign gives direction, size gives tightness of the linear pattern. r = 0.8 is a strong march together; r = 0.2 a loose drift. Correlation is symmetric and unit-free, which makes it comparable — and makes it silent about cause in both directions.

Scatterplots first

Plot before you compute. A scatterplot reveals curves, clusters and outliers that a single coefficient flattens: a strong curve can hide behind r near zero, and one wild point can manufacture a correlation from noise. The eye is the first regression diagnostic.

Simple regression

Regression fits the line that best predicts one variable from another and reports its slope in real units — each extra hour of study associated with so many extra marks. The slope's interval and the scatter around the line matter more than the line's elegance; prediction without its error

band is fortune-telling.

The causal warning

Association earns causal language only through design: random assignment, or careful control of confounders with honest argument. Fertiliser use correlates with yield partly because richer farmers buy both fertiliser and better land. Name the confounders you fear in every observational claim you publish.

Key points

- Plot first; compute second.
- Slopes speak in real units — report them with intervals.
- Association plus design earns causal words; association alone does not.

Exercises

1. Sketch scatterplots for r ≈ 0.9, r ≈ 0, and a strong curve with r ≈ 0.

2. Interpret: 'each extra hour of study, +1.8 marks (interval 1.2 to 2.4)'.

3. Name two confounders in the fertiliser–yield association.

Chapter 6

Designing Studies That Can Conclude

The final chapter joins analysis to design: randomisation, controls, and the reporting standards that let a study's conclusions be believed — the craft that separates research from ritual.

Random assignment

Assigning treatments by lot breaks the link between treatment and every confounder, known and unknown. It is the single most powerful sentence in a methods section. Where randomisation is impossible, say so, and argue your comparisons with extra care rather than borrowed confidence.

Controls and comparisons

A control group shows what would have happened anyway — the season's rain, the market's rise. Without it, every improvement is credited to your intervention and every decline blamed on it. The history of failed projects is largely the history of missing controls.

Pre-commitment and honest reporting

Write the question, hypotheses, sampling and analysis plan before collecting data, and report what you actually did — including the outcomes that disappointed. Selective reporting converts the scientific record into advertising, one omitted result at a time.

Your road as an analyst

Master one good software tool, keep your code and data organised so another researcher can rerun your analysis, and seek review before you publish. Statistics is a craft of honesty enforced by technique; this volume has given you the technique — the honesty is yours to keep.

Key points

- Randomisation defeats confounders, known and unknown.
- No control group, no causal credit.
- Plan before data; report everything, including disappointments.

Exercises

1. Design a randomised trial of a study-skills course for 60 students.

2. List what a control group protects against in a fertiliser demonstration.

3. Write the five headings of a pre-committed analysis plan.

Appendix — Worked Examples and Practice

From Chapter 1 — Sampling and the Sampling Distribution

Worked example: with standard deviation 40 and n = 100, the standard error of the mean is 40/√100 = 4. To halve it to 2 you need n = 400, because √400 = 20 and 40/20 = 2. The square-root law in action: each halving of error costs four times the fieldwork.

From Chapter 2 — Estimation and Confidence Intervals

Worked example: 46/200 = 0.23. The standard error of a proportion is √(p(1−p)/n) = √(0.23×0.77/200) ≈ 0.030, so the 95 per cent interval is roughly 0.23 ± 1.96×0.030 ≈ 0.23 ± 0.06 — between 17 and 29 per cent. The district's true rate is plausibly anywhere in that band; policy should be written for the band, not the point.

From Chapter 3 — Significance Tests — Logic and Limits

Worked example: two teaching methods differ by 0.4 marks in 10,000 students, p = 0.01. Statistically surprising, educationally nothing — less than half a mark. Contrast a drug trial of 30 patients where survival rises 20 points at p = 0.08: not 'significant', possibly life-saving, and demanding a bigger trial rather than dismissal. Size and surprise are different questions.

From Chapter 4 — Comparing Groups — t-tests and Chi-square

Worked example: adoption of composting — Village A: 34 of 50 households adopt; Village B: 22 of 50. The table is [[34,16],[22,28]]. Chi-square ≈ 5.8, p ≈ 0.016: the difference (68% vs 44%) is larger than chance commonly produces at these sizes. Report table, percentages, test and p together — each checks the others.

From Chapter 5 — Relationships — Correlation and Regression

Worked example: study hours and marks for ten students give r = 0.82 and a slope of +1.8 marks per hour. The honest report: a strong association in this class, worth about two marks per hour across the observed range — and no proof that forcing hours would force marks, since keen students both study more and learn faster per hour.

From Chapter 6 — Designing Studies That Can Conclude

Worked example: sixty students volunteer for a study-skills course; thirty places exist. Drawing lots for the places is both the fairest allocation and a randomised trial: at term's end, the mark difference between the drawn and the undrawn estimates the course's effect cleanly, because keenness — the great confounder — was distributed by the same lot.

Further Practice — Review and Discussion

Work through these prompts in writing; they are the revision spine of this book. Answer in full sentences, with examples from your own place of study wherever the prompt allows.

Chapter 1 — Sampling and the Sampling Distribution

- Explain in your own words, with one example: Sampling variation is predictable, not shameful.
- Explain in your own words, with one example: Standard error shrinks with the square root of sample size.
- Explain in your own words, with one example: Inference cannot repair a biased sample.

Discussion: where does the subject of this chapter — sampling and the sampling distribution — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Chapter 2 — Estimation and Confidence Intervals

- Explain in your own words, with one example: Report the interval, not only the point.
- Explain in your own words, with one example: 95 per cent describes the method's long-run success.
- Explain in your own words, with one example: Interval width honestly reports your uncertainty.

Discussion: where does the subject of this chapter — estimation and confidence intervals — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Chapter 3 — Significance Tests — Logic and Limits

- Explain in your own words, with one example: Tests measure surprise under the null — they never prove the alternative.
- Explain in your own words, with one example: p is not the probability that the null is true.
- Explain in your own words, with one example: Report effect size and interval with every p-value.

Discussion: where does the subject of this chapter — significance tests — logic and limits — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Chapter 4 — Comparing Groups — t-tests and Chi-square

- Explain in your own words, with one example: t-test for means; chi-square for counts in categories.
- Explain in your own words, with one example: Paired designs are stronger: each unit is its own control.
- Explain in your own words, with one example: Always print the table or the group summaries beside the test.

Discussion: where does the subject of this chapter — comparing groups — t-tests and chi-square — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Chapter 5 — Relationships — Correlation and Regression

- Explain in your own words, with one example: Plot first; compute second.
- Explain in your own words, with one example: Slopes speak in real units — report them with intervals.
- Explain in your own words, with one example: Association plus design earns causal words; association alone does not.

Discussion: where does the subject of this chapter — relationships — correlation and regression — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Chapter 6 — Designing Studies That Can Conclude

- Explain in your own words, with one example: Randomisation defeats confounders, known and unknown.
- Explain in your own words, with one example: No control group, no causal credit.
- Explain in your own words, with one example: Plan before data; report everything, including

disappointments.

Discussion: where does the subject of this chapter — designing studies that can conclude — touch your own community, farm, school or workplace? Write one page, naming real places and real constraints, and end with one question you would research next.

Glossary of Key Terms

Standard error — The standard deviation of a sample statistic over repeated samples.

Confidence interval — An interval built by a method that captures the true value at a stated long-run rate.

Null hypothesis — The assumption of no effect, against which data are judged.

p-value — The probability, under the null, of results at least as extreme as those observed.

t-test — A test comparing means relative to within-group scatter.

Chi-square test — A test comparing counts across categories.

Regression — Fitting the relationship that predicts one variable from another.

Confounder — A variable linked to both treatment and outcome that can fake or hide an effect.

End of book. Nji Griphan · griphan.48hourrfprescue.com · info@griphan.48hourrfprescue.com