

Analysis Studio: The Complete How-To Guide

Setup, survey designs, licensing, and all twenty modules, step by step.

MAINEBLOCK · NUMBERS, DONE RIGHT.

October 2026 · maineblock.com

Contents

1. Opening Analysis Studio
2. A tour of the screen
3. Getting your data in
4. Your first run
5. Free and subscription: what the difference is on screen
6. Buying a subscription and entering your key
7. Variables and roles
8. Survey designs and weights
9. Reading results, copying, and downloading
10. Module 1: Descriptive statistics and frequencies
11. Module 2: Crosstabs and tests of association
12. Module 3: Comparing means: t-tests and one-way ANOVA
13. Module 4: Correlation
14. Module 5: Linear regression
15. Module 6: Logistic regression
16. Module 7: Sample size and power planner
17. Module 8: Rank-Based Tests and Exact Alternatives
18. Module 9: Key driver analysis
19. Module 10: A/B test analyzer
20. Module 11: Survey weighting and raking
21. Module 11, second procedure: Propensity adjustment for nonprobability samples
22. Module 12: Trend analysis across survey waves
23. Module 13: Scale reliability and exploratory factor analysis
24. Module 14: Count, ordinal, and multinomial regression
25. Module 15: Mixed-effects models and repeated measures
26. Module 16: Survival analysis
27. Module 17: Time series and forecasting
28. Module 18: Propensity score analysis
29. Module 19: Statistical Quality Control and Improvement
30. Module 20: Epidemiologic Analysis
31. Saving and sharing your work
32. Troubleshooting
33. Glossary

Opening Analysis Studio

1. Open your browser and go to **app.maineblock.com/analysis-studio.**
2. Wait for the page to finish loading. You should see the heading "Analysis Studio" and the line "Advanced analytics, right in your browser."
3. Look at the top right of the header. You should see two controls, "Appearance" and "Theme". They change how the app looks, never what it computes. "Appearance" offers "Light", "Dark", and "System". "Theme" offers nine accent colors, starting with "Signal Orange".
4. Scroll down. The first card is a bright orange card headed "Start Your Free Trial!", with the lines "Run our models free until March 31st, 2027" and "Downloads and copying need a license key.". Below it you should see three cards in order, each headed by an icon tile and a plain name: "Data", "Survey design", and "Analysis".

Nothing needs an account, and nothing needs installing. The statistics engines are native code that runs in the page itself, so the app is ready as soon as the cards appear.

> Your data stays on your machine. Analysis Studio makes exactly one kind of network call: license verification, and it sends the license key alone. No dataset, result, or metadata is uploaded, ever. Every number in this guide is produced on the reader's own device.

A tour of the screen

The screen is a single page of stacked cards. From top to bottom:

- **"Start Your Free Trial!"** is the orange card at the top of the page. It is an announcement only; it holds no controls. Under its heading, two lines read "Run our models free until March 31st, 2027" and "Downloads and copying need a license key." The license key lives in a pop card that opens only when you attempt a download or "Copy Results" without one (or when you press the PrintScreen key). The section "Free and subscription: what the difference is on screen" explains what the key changes.
- **"Data"** is where a dataset enters. It holds a file picker labeled "Data file (CSV, TSV, SPSS, Stata, SAS, or Excel)" and a button, "Load the sample data". Under them, a profile of the loaded data appears: the row count, the variable count, and a table with one row per variable showing its name, its level, and its valid, missing, and distinct counts. Level means the measurement level the app assigned the variable, such as continuous, count, nominal, ordinal, or binary. The level decides which modules and which roles a variable can fill.
- **"Survey design"** declares how the data were sampled: weights, strata, clusters, or replicate weights. Most first runs leave it alone, and Part 2, in the section "Survey designs and weights", covers it in full.
- **"Analysis"** is where the work happens. It starts with the line "Choose a module to set it up and run it.", then the module picker: a "Find a module" search box with a count beside it (the roster holds twenty modules and twenty-one procedures, because Module 11 carries two), then the modules as cards grouped under seven headings, "Describe", "Compare and model", "Plan and test", "Survey methods", "Advanced methods", "Quality and improvement", and "Epidemiology", each heading shown as a dark band in capital letters. Each card carries an icon, "Module N" and the module's name, a pill reading "Free" or "Free trial", and a one-line description. Typing in the search box narrows the cards to the ones whose name or description matches, and clearing it brings every card back; searching never selects a module or changes one. The "Survey methods" group holds a twenty-first card, "Propensity adjustment for nonprobability samples", labeled "Module 11 · Second procedure"; it is the second procedure of Module 11 and has its own section in that chapter. Choosing a module opens its setup in a panel below the picker, and the app brings that panel to the top of the screen for you, so there is nothing to scroll to or hunt for. The setup is headed by the module's icon, its "Module N" label, its name, and its description. Its fields are grouped under titled sections: the variable pickers first, under "Variables", each shown as a well with a dashed border and a note of how many variables to choose, then an "Options" section, then a "Random seed" field preset to 42, then the "Run" button. The picker stays above the setup with the chosen module still marked. Module 1 is the first module on the newer step layout: numbered steps instead of the "Variables" and "Options" sections,

variable pickers that hold the chosen variables as chips, a "You will get" card at the side, and an action bar with "Run analysis" at the bottom of the panel. Its chapter describes that screen in full.

- **Results** appear inside the module's own panel, below the setup: first a "Checks" list when the module has anything to flag, then the result tables, then a row of download buttons, and a "Run manifest" summary you can open to see exactly what was run.

Two habits pay off everywhere. First, the variable pickers in a module form fill from your loaded dataset and show each variable with its level in parentheses, so load data before choosing a module. Second, every module form shows a hint under each picker telling you how many variables that role takes and which levels it accepts. Read the hint before you pick; it is the module's own contract.

Getting your data in

Analysis Studio reads delimited text files, CSV (comma-separated) and TSV (tab-separated), and the standard statistical formats: SPSS (.sav), Stata (.dta), SAS (.sas7bdat, and the SAS transport format .xpt), and Excel workbooks (.xlsx, first sheet). One layout rule holds for every format: one row per case and one column per variable, with the variable names in the first row of a text file, or carried in the file itself for the statistical formats. One refusal to know: compressed SPSS files (.zsav) are not in the supported list, and the app says so on screen: "Compressed SPSS files (.zsav) are not in the supported format list. Save the file as a standard .sav and try again."

1. Prepare your file so the first row holds the variable names and every following row is one case, in the same column order throughout.
2. In the "Data" card, choose the file with the picker labeled "Data file (CSV, TSV, SPSS, Stata, SAS, or Excel)".
3. You should see the profile appear under the card: "N rows, M variables", then the variable table.
4. Read the variable table before running anything. Check the row count against your file, check each variable's level (a variable you meant as a category that arrives as continuous will not appear in category pickers), and check the missing counts so no column silently failed to read.

If you have no file handy, press "Load the sample data". The app loads a twelve-row survey sample with four variables: `y` (a continuous outcome), `w` (a sampling weight), `stratum`, and `psu` (the sampled cluster). The next section runs a first analysis on it.

Three facts about data handling are worth knowing now. Missing values are cells the file leaves empty or marks in the usual ways; modules exclude a case from a statistic when a variable that statistic needs is missing, and they say so in their notes. A variable the app reads as an identifier (a case ID, in effect) is excluded from statistics by Module 1 and flagged, because an ID is a label, not a measurement. And nothing about your file leaves the browser: the profile, the parsing, and every run happen on your device.

Your first run

This walkthrough runs Module 1, "Descriptive statistics and frequencies", on the sample data, from an empty page to a finished result. Every later module chapter assumes you have done this once.

1. Press "Load the sample data" in the "Data" card. You should see "12 rows, 4 variables" and the variable table listing `y`, `w`, `stratum`, and `psu`.
2. Leave the "Survey design" card as it stands for this run. Its preset is "None (simple random sample)", so no design is applied and the statistics are the plain, unweighted kind. The section "Survey designs and weights" explains how to declare a design when your data have one.
3. In the "Analysis" card, choose "Module 1: Descriptive statistics and frequencies" in the module picker. The setup opens below the picker and the app brings it to the top of the screen. It is headed by the module's icon, the label "Module 1", its name and description, a "Free" pill, and a link, "How this module works". Below the head stand three numbered steps, a side panel with "You will get" and "Survey design" cards, and an action bar at the bottom.
4. In Step 1, "Choose what to summarize", click in the "Variables to summarize" box and choose `y` from the list. It appears as a chip tagged "number". The line under the box reads "Numbers get averages and spread. Categories get frequency tables." Any variable can be described; `y` is the continuous one.
5. Leave "Compare by (optional)" empty. Leave Step 2 on the "Essentials" preset, leave Step 3 as it stands, and leave "Advanced settings" closed; its summary line reads "Wilson intervals · Logit under a survey design · Quantile type 7 (R default)", and the "Random seed" inside it stays at 42. The action bar at the bottom now reads "1 variable · Essentials preset".
6. Press "Run analysis" in the action bar. You should see "Running…" briefly, then the results.

Read the result top to bottom. Above the first table the app prints a short takeaway sentence: "The average y was 34.58 (95% CI 23.72 to 45.45; n = 12)". The first table is titled "Descriptive statistics: y". On the sample's twelve values (10, 12, 20, 22, 24, 30, 40, 44, 48, 50, 55, 60) it reads: n 12, Missing 0, Mean 34.58, SD 17.10, SE of mean 4.94, CI lower 23.72, CI upper 45.45, Median 35, Q1 21.5, Q3 48.5, Minimum 10, Maximum 60.

The table prints full precision (the mean shows as 34.5833...), so read the figures here and throughout this guide as the rounded values the takeaway uses. After the tables come a methods paragraph stating exactly how the figures were computed, and the row of download buttons.

Setup is complete when that table is on your screen. You have loaded data, chosen a module, set a form, run it, and read a result: the five moves every module in this guide repeats.

Free and subscription: what the difference is on screen

Analysis Studio costs $299 per year for the subscription, one paid level. What the subscription buys is specific, and the screen shows you the line wherever you meet it.

The promotion, first, because it is what you will see today. Every module is free to use, meaning free to run on screen, through **March 31, 2027.** During the promotion you can open and run all twenty modules, including every advanced one, with no key at all. Two things stay behind the subscription during the promotion, exactly as they are after it: every download (the CSV, the Word and Excel reports, the R script, the SPSS syntax, and the run manifest) and the "Copy Results" button. The subscription today buys the downloads and copying, and from April 1, 2027 it also buys running the advanced modules.

The standing arrangement, from April 1, 2027. Nine modules are free to run on screen for anyone: Module 1 (Descriptive statistics and frequencies), Module 2 (Crosstabs and tests of association), Module 3 (Comparing means: t-tests and one-way ANOVA), Module 4 (Correlation), Module 5 (Linear regression), Module 6 (Logistic regression), Module 7 (Sample size and power planner), Module 11 (Survey weighting and raking), and Module 12 (Trend analysis across survey waves).

Twelve procedures need the subscription to run: Modules 8, 9, 10, 13, 14, 15, 16, 17, 18, 19, and 20, plus the second procedure inside Module 11, "Propensity adjustment for nonprobability samples". Downloads and copying need the subscription on every module, free ones included. That split is the permanent arrangement the module chapters are written against: each chapter's access line gives the standing rule and notes the promotion.

How the line feels in use:

1. Choose any module during the promotion. Its form opens and "Run" is active.
2. After the promotion, choose a subscription module without a verified key. The form opens (you can read every control), but "Run" is disabled and a note beside it reads "This module requires a license."
3. Run a free module without a key, then press any download button or "Copy Results". A pop card titled "License Key Required" opens instead of the download: downloading and copying wait there while you enter a key or buy one. Nothing about the run itself is interrupted, and running modules never opens the card.
4. In the pop card, with no key entered yet, there is no status line under the buttons at all: the card simply waits. The status line appears only when a key is in play: "Verifying…" while the check runs, then the outcome, as the next section describes.

One thing is never behind the line: correctness. Applying weights you already have, with design-correct standard errors, inside a free module is ordinary platform behavior and is never gated. The free and subscription paths run the same engines through the same checks.

Buying a subscription and entering your key

1. The purchase path is the "Buy a License Key" button on the pop card (the card that opens when you attempt a download or "Copy Results" without a key), or the "Buy" path for Analysis Studio on the product page at maineblock.com. Both hand you to the checkout on this app's own address. Checkout is handled by Stripe; the price is $299 for one year.
2. When the purchase completes, a license key arrives by email. The key has the shape AS-XXXX-XXXX-XXXX-XXXX: four groups of four characters after the AS prefix.
3. Back in Analysis Studio, attempt the download or copy you wanted. The pop card opens; type or paste the key into its "License key" field.
4. Press "Verify key". You should see "Verifying…" and then "License verified." with the line "All modules and outputs are unlocked." and, when the server provides one, an expiry date. The card closes and the download or copy you attempted finishes on its own.
5. From then on, downloads and "Copy Results" work directly, with no card, for as long as the key stays verified.

The app remembers the key in your browser and re-verifies it on your next visit, so you enter it once per browser. "Forgot key", on the pop card, opens the key recovery form: "Lost your key? Enter the email you bought with. A new key will be sent to it, and the old key will stop working." Enter the email address the purchase was made with and press "Send a new key". The form answers "If that email has an Analysis Studio purchase, a new license key is on its way to it." The answer is the same whether or not the address has a purchase, so the form cannot be used to find out who bought.

The replacement key arrives by email. It belongs to the same purchase and keeps the same expiry date, and from that moment it is the only live key for the purchase: the key in the original purchase email stops working. The form also carries a "Clear the key saved in this browser" button; that button is what removes the saved key and returns the app to the unlicensed state.

To reach the card when no download is pending, press the PrintScreen key (on keyboards that have one) or simply attempt any download. "Manage", the subscription management path linked from your purchase and renewal emails, is where renewal and cancellation live; cancelling turns off auto-renewal while the license stays valid through the paid year. If a key does not take, the status line in the pop card says why in the server's own words (an unrecognized key, an expired one), and the fix is to re-enter the key exactly as it arrived, hyphens included, before anything else.

Two take-away paths deserve a plain word. Printing results is not offered: a printout of the page shows a short note where the results would be, because the printable copies are the Word and Excel report downloads, which are part of the licensed set. And while no key is verified, pressing the PrintScreen key puts a short license message on the clipboard in place of the screen capture and opens the pop card. No web app can stop every capture (a photograph of the screen, for one); these are the measures the app takes on the paths a web page can reach.

A practical note for the promotion period: a key is worth entering as soon as you have one. It unlocks downloads and copying immediately, and it carries you through April 1, 2027 with nothing further to do.

Variables and roles

Every module form is built from roles. A role is a slot with a job: "variable" for the things to describe, "outcome" for the thing to explain, "predictors" for the things that explain it, "group" for the split to compare across, and so on. Each chapter in Part 3 lists its module's roles; this section covers the rules they all share.

- A picker only offers variables whose level fits the role. A "group" picker accepts nominal, ordinal, or binary variables; a variable you know is a category but the app read as continuous will not appear until the data itself is fixed (recode it in your file so its values are category labels or codes the parser reads as categories, then reload).
- The hint under each picker states the count range and the accepted levels, for example "Choose 1 to 30" for a predictor list. **Minimum counts are real requirements:** a module whose form validates will stop under the heading "Check the setup" and name the missing piece, rather than guessing.
- Levels come from the file itself. Text columns arrive as text or nominal; columns of whole numbers arrive as count; columns with exactly two distinct values arrive as binary; ordered labels arrive as ordinal when the file marks them so.
- The "by" and "group" roles take category variables; the "predictors" roles in the regression modules take any measured level, because a categorical predictor is expanded into indicator columns automatically, with a reference level named in the output.
- "Random seed", preset to 42, sits under every form. Procedures that simulate or bootstrap (Monte Carlo tests, bootstrap intervals, matching) use it, so the same data, the same settings, and the same seed give the same result. Change it when you want a different simulation draw; record it when a number must be reproducible later.

> A note on naming. In this guide, a control is always called by the label the screen prints. The module forms label their option controls in plain words (for example "Standard-error type" in the regression modules); each chapter spells out what each one does in plain words next to its label.

Survey designs and weights

If your data come from a survey, the sample design changes the standard errors, and the "Survey design" card is where you declare it. Nothing here is required for data you treat as a simple random sample.

1. In the "Design" list, choose how the sample was drawn: "None (simple random sample)", "Weights only", "Taylor (strata and clusters)", or "Replicate weights". The preset is "None (simple random sample)": no design is applied unless you declare one here, and the "Weight column", "Strata column", and "Cluster (PSU) column" lists start unchosen, each showing "(none)", until you pick a column yourself. Nothing is filled in for you.
2. For "Weights only" or "Taylor", set "Weight column" to your weight variable, and "Weight type" to "sampling", "frequency", or "analytic". Sampling weights are the survey default: each case stands for that many population cases. Frequency weights say each row repeats that many times, and are rescaled before use. Analytic weights say the case values are means built on that many observations.
3. For "Taylor", also set "Strata column" and "Cluster (PSU) column" when your design has them, and "Population PSU count column (FPC)" when you have population cluster counts for a finite-population correction. Set "Lonely PSU" (the rule for a stratum containing a single cluster) to "adjust", "average", "remove", or "certainty".
4. For "Replicate weights", set "Replicate weight columns", "Replicate type" ("JK1", "JKn", "BRR", "Fay", or "bootstrap"), "Replicate scale (other)" when the type is "other", and "Fay rho" (preset 0.5) for the Fay type.
5. Run any module. Under a declared design, the modules compute design-based standard errors (Taylor linearization or the replicate method) and print in their methods paragraph which design was in force.

Two cautions. Declaring a design you do not have invents precision or destroys it; declare only what the data collection actually did. And scale matters for frequency and analytic weights in a way it does not for sampling weights, so when in doubt, read the note Module 1 prints about the weight type you declared: the modules carry a warning for each pairing of weight type and method, and they print it on the result rather than hiding it.

Reading results, copying, and downloading

Every run ends in the same place: a result block inside the module's panel. Its anatomy is the same in all twenty modules.

- A **"Checks"** list comes first when the module has something to flag: a design feature that changed the method, a sparse cell, a fitted-model warning. Read checks before tables; they are the module telling you what it did about your data.
- The **result tables** follow, titled by what they hold ("Group statistics", "Model fit", "Driver importance", and so on). Each chapter in Part 3 walks its module's tables in order. A takeaway sentence in words is printed above the first table, and a methods paragraph records the exact procedure, options, and design in force, so the result can be defended later. The tables print full precision (long decimals); the takeaway sentence rounds to two decimals, and this guide quotes the rounded figures.
- The **download row** offers "Download CSV", "Copy Results", "Download Word Report", "Download Excel Report", "Download R Script", "Download SPSS Syntax", and "Download Run Manifest". CSV and the two scripts cover the latest run; the report files and the manifest cover the session report, which accumulates every run since the page loaded (the note under the buttons keeps count of the result sets so far). The R script and SPSS syntax reproduce the analysis in those environments where the module supports it.
- The **"Run manifest"** summary opens a full record of the run: module, inputs, options, seed, design, and the checks that fired. When a result is questioned, the manifest is the answer.

Downloads and "Copy Results" need the subscription at all times (see "Free and subscription: what the difference is on screen"). Without a verified key, pressing one opens the "License Key Required" pop card instead of downloading; with one, the action runs directly and "Copy Results" confirms "Results copied to the clipboard." on success. One more failure mode is the browser's, not the app's: if the browser refuses clipboard access, the app reports "Copy failed: the browser refused clipboard access." and the fix is to allow clipboard permission for the page, or to download instead.

Module 1: Descriptive statistics and frequencies

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

For continuous variables, a full descriptive table: count, missing count, mean, standard deviation, standard error of the mean, a confidence interval for the mean, median, quartiles, minimum, and maximum, with more statistics on request. For categorical variables, a frequency table: counts, percents, valid percents, and a confidence interval for each category's share. With a grouping variable in "Compare by", the same tables split by group. Under a declared survey design, the figures are the design-based counterparts (weighted means, proportions, and linearization or replicate standard errors), and the output says so.

The data it needs

One or more variables of any kind in the "Variables to summarize" picker, and up to three category variables in "Compare by" if you want the tables split by group. A variable that cannot group, such as a number, still appears in the "Compare by" list, shown dimmed with the reason stated under its name. The built-in sample data works for a first pass (see "Your first run"). For frequencies, use a small categorical set: a variable `region` with twelve cases reading North, North, South, South, South, West, West, West, West, East, East, North, which gives counts of 3, 3, 4, and 2 across the four categories, or 25.0, 25.0, 33.3, and 16.7 percent.

The steps

1. Load data, and declare the design in "Survey design" if you have one; the walkthrough in "Your first run" shows the plain path. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card.
2. Choose "Module 1: Descriptive statistics and frequencies". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: it counts the summary tables, the frequency tables, the statistics per number, and the charts the run will produce.

3. Step 1, "Choose what to summarize": pick the variables in "Variables to summarize". Each chosen variable becomes a chip with a tag: "number", "category", "count", or "date". Add up to three grouping variables in "Compare by (optional)" if you want split tables; its chips carry the group count (such as "4 groups") instead of a tag. The line under the first box reads "Numbers get averages and spread. Categories get frequency tables."; the line under the second reads "Categories only. Results are shown for each group. Up to 3 grouping variables."

4. Step 2, "Choose what to report": start from the "Preset" control. "Essentials" is the shipped set: n, Missing, Mean, SD, SE of mean, the confidence interval, Median, Q1, Q3, Minimum, and Maximum. "Detailed" switches every statistic on and adds the 25th, 50th, and 75th percentiles. Switching any statistic by hand moves the preset to "Custom". The statistics sit in five groups:

- "Typical value": "Mean", "Median", "Mode", and "Trimmed mean". - "Spread": "Standard deviation", "Variance", "Minimum and maximum", and "Range". - "Quartiles": "First and third quartiles" and "Interquartile range". - "Precision": "Confidence interval" and "Standard error", with the "Level" beside them: 90%, 95%, or 99%, printed on the output as "90% CI", "95% CI", or "99% CI". - "Counts and shape": "Valid n and missing", "Skewness", "Kurtosis", and "Total". Skewness and kurtosis each print with their standard error. - "Percentiles": extra percentiles, added one at a time as whole numbers from 1 to 99 with "+ Add percentile". Each appears as a chip, such as "25th", and as a column named like P25 in the table.

5. Step 3, "Frequency tables": three switches for the category tables:

- "Show categories with no cases": on by preset, so labeled categories with no cases still print, with a zero count. The line under it reads "Keeps labeled but empty categories in the table." - "Cumulative percent": adds a "Cumulative valid percent" column. The line under it reads "Useful for ordered categories." - "Weighted counts": adds a "Weighted count" column when a survey weight is set. The line under it reads "Adds a column when a survey weight is set."

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Wilson intervals · Logit under a survey design · Quantile type 7 (R default)". Inside:

- "Proportion interval": "Wilson score" (preset) or "Clopper-Pearson (exact)" for category-share intervals without a design. - "Proportion interval under a survey design": "Logit" (preset) or "Korn-Graubard" for category-share intervals under a design. - "Quantile definition": "Type 7 (R default)" (preset, the usual linear-interpolation quartiles) or "Type 6 (SPSS)", the definition SPSS uses, without a design. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7.

Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work (for example, "3 variables · grouped by region · Essentials preset"), or names the first thing still missing ("Choose a variable to run"). "Reset", in the same bar, restores this module's presets, and asks first when anything was changed. You should see the checks list (if any), then one table per variable, each titled "Descriptive statistics: <variable>" or "Frequencies: <variable>".

Reading the output

The continuous table's first column is headed "Variable", or "Group" when a "Compare by" variable splits it; each row is one statistic. The frequency table's columns are Group, Category, Count, Percent, Valid percent, CI lower, and CI upper, with the optional columns appearing when their options are on. Percent uses every case as its base; valid percent excludes missing cases, and the two differ exactly when the variable has missing values.

A note flags any category whose share's interval was computed under a design, and identifier variables (case IDs) are excluded with a note, since a mean ID is not a statistic. Above the main table, the takeaway sentence states the headline in words, and the methods paragraph after the tables records the definitions used, including which interval and which quantile definition were in force.

Under a design, two printed behaviors change. The mode and the trimmed mean are not computed under a design (the output notes this), because their design-based forms are not standard; and the design note sentences name the weight type and the variance method so the table can be quoted safely.

Downloads and copying

All seven outputs are available with a subscription: CSV of the tables, "Copy Results", the Word and Excel reports, the R script, the SPSS syntax, and the run manifest. Without one, the buttons show the licensed-set note.

When it refuses or warns

- With no variable selected, the action bar reads "Choose a variable to run"; a run attempted anyway stops under "Check the setup" with "Select at least one variable."
- A fourth grouping variable cannot be added: at three, the "Compare by" list shows the remaining variables dimmed, with the limit stated.
- A percentile is refused at the box unless it is a whole number from 1 to 99: "Enter a whole percentile between 1 and 99." A duplicate is refused with "That percentile is already in the list."
- Expect warnings, not stops, for design and weights pairings the module handles with a printed convention (for example, how a replicate design treats proportion intervals), and for cumulative percents requested on a variable where they do not apply.

Module 2: Crosstabs and tests of association

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

Cross-tabulations of two categorical variables, with counts and percentages and a battery of association tests: Pearson and likelihood-ratio chi-square, Fisher exact tests (including a Monte Carlo form for larger tables), Cramér's V (with a bias-corrected form), and for 2 by 2 tables the odds ratio, risk ratio, and risk difference. Layered tables either print separately per layer or combine through the Cochran-Mantel-Haenszel test. A single categorical variable gives a goodness-of-fit test against equal or stated proportions. Under a design, the tables carry weighted counts and design-based (Rao-Scott corrected) tests.

The data it needs

One "row" variable and one "column" variable, both categorical, and an optional "layer" variable. Worked example, forty respondents cross-classified by group and answer: group A answers Yes 14, No 6; group B answers Yes 8, No 12. Within its own group, Yes is 70.0 percent of A's twenty answers and 40.0 percent of B's twenty.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 2: Crosstabs and tests of association". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it counts the cross-table, names the percentages that will print beside the counts, lists a chi-square test of association with effect sizes, and, when a layer variable is chosen, adds either a separate table for each layer level or one combined test across the layers.

3. Step 1, "Choose variables": the step description reads "Pick the row and column variables for the table. Add a layer variable to repeat the table within each of its groups."

- "Row variable": click in the box and choose one category variable. It appears as a chip; grouping roles carry the group count, such as "4 groups", instead of a type tag. The line under the box reads "One category variable. Its levels run down the side of the table." - "Column variable": choose one category variable in the same way. The line under the box reads "One category variable. Its levels run across the top of the table." - "Layer variable" (optional): choose one category variable to repeat the table within each of its groups. The line under the box reads "Optional. A separate table is printed for each

level, with a combined test." - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "spend is a number; this role needs a category or a two-level category."

4. Step 2, "Set up the table": the step description reads "Choose what prints in the cells and which tests and measures run." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" shows total percentages only, with every test and measure switch off. "Detailed" shows row, column, and total percentages and switches every test and measure on. Both presets set the confidence level at 95%. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in three groups:

- "Table contents": "Percentages shown", a set of three boxes, "Row", "Column", and "Total". The line under the group reads "Which percentages print beside the count in each cell." - "Tests and measures": four switches. "Bias-corrected effect size" reports the bias-corrected form of the association measure for the table. "Risk difference" gives, for 2 by 2 tables, the difference between the two row proportions, with an interval; the interval is the Newcombe score interval. "Column proportion tests" tests which columns differ from each other, pair by pair, with an adjustment for the number of tests; the adjustment is Holm, applied within rows. "Adjusted standardized residuals" shows which cells sit far from what independence would predict. - "Precision": "Confidence level", the level of the intervals printed with the table measures: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Yates continuity correction: Off · One-way table grouping: Not set · Layer handling: Separate table per layer · Exact test for larger tables: Not run · Simulated tables: 10000 · Random seed: 42". Inside:

- "Yates continuity correction": a switch. The line under it reads "Applies to 2 by 2 tables. Makes the chi-square test more conservative with small counts." - "One-way table grouping": "Row variable" or "Column variable". The line under it reads "For a table with a single variable, which variable its summary is grouped by." - "Layer handling": "Separate table per layer" or "Cochran-Mantel-Haenszel (CMH)". The line under it reads "Print a separate table per layer, or combine the layers into one adjusted test." - "Exact test for larger tables": "Not run" or "Monte Carlo simulation". The line under it reads "For tables larger than 2 by 2, estimate the exact test by simulation instead of the chi-square approximation." - "Simulated tables": how many tables the simulation draws for the exact test; the preset is 10000. The line under it reads "More tables give a steadier estimate." - "Expected proportions": a chip input for a one-way table. Add one proportion at a time; each appears as a removable chip. The line under it reads "For a one-way table, the share expected in each category, in category order. Leave empty to expect equal shares." Proportions that are not positive, or that do not sum to one, stop

a goodness-of-fit run with the reason printed. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "region by answer · Essentials preset", or names the first thing still missing: "Choose a row variable and a column variable to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the crosstab titled "Crosstab: <row> by <column>", then "Tests of association" and "Effect sizes and 2 x 2 measures".

Reading the output

The worked example prints the crosstab with counts and total percents, each cell as a share of all forty respondents (A Yes 14 at 35 percent, A No 6 at 15 percent, B Yes 8 at 20 percent, B No 12 at 30 percent), a Pearson chi-square of 3.64 on 1 degree of freedom ($p = 0.057$), and Cramér's V of 0.30. (The chi-square's expected counts, 11.0 and 9.0 in each row, are computed for the test and its sparseness checks; they are not printed as a table.) Read the tests table first for the verdict, then the effect-size table for the size of the association, then the residuals (when on) for which cells carry it.

For a 2 by 2 table, a note under the effect measures states their orientation: the odds ratio and relative risk compare the row variable's first level with its second, for the outcome named in the effect table's title, and the risk ratio is the share with that outcome in the first row divided by the share in the second; phi is signed on the same orientation. When the outcome takes the values 0 and 1, the measures describe level 1, the event, by default; for any other pair of outcome levels they describe the column variable's first level.

Under a design, the table title gains "(design-based)" and the tests are the Rao-Scott forms. The goodness-of-fit variant prints its own table when only one variable is chosen.

Downloads and copying

The full set of seven outputs, subscription permitting, including an R script and SPSS syntax that rebuild the same tables and tests.

When it refuses or warns

- Expected proportions that are not positive, or that do not sum to 1 for a goodness-of-fit run, stop the run with the reason printed.
- A continuous variable in a categorical role stops the run: "crosstabs need categorical variables", naming the variable and its level.
- Sparse tables produce checks rather than stops: small expected counts are flagged so the exact tests, not the chi-square approximation, carry the conclusion.

Module 3: Comparing means: t-tests and one-way ANOVA

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

Four mean comparisons in one module: a one-sample t test against a stated value, an independent two-sample test in Welch's form (unequal variances allowed), a paired t test on the difference of two variables, and a one-way ANOVA across the groups of a categorical variable, with post hoc pairwise comparisons and effect sizes. The procedure can be chosen or left on "Automatic", which reads the structure you gave it. Under a design, the counterparts are design-based: the two-group test in the survey t-test form, the multi-group test as a design-based Wald test, and paired and one-sample tests on design-based differences.

The data it needs

One continuous "outcome". Then either a categorical "group" (two groups for the t test, more for ANOVA), or a second continuous "paired" variable, or neither (one-sample). Worked example: scores for group A of 22, 25, 28, 24, 26 (mean 25.0) and group B of 18, 20, 19, 22, 21 (mean 20.0). ANOVA example: three groups of four; group A scores 5, 6, 7, 6 (mean 6.0), group B scores 8, 9, 7, 8 (mean 8.0), and group C scores 11, 12, 10, 13 (mean 11.5).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".
2. Choose "Module 3: Comparing means: t-tests and one-way ANOVA". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it names the comparison the chosen variables call for (a one-sample test, a paired comparison, or a comparison across groups), adds group means with the difference and its interval, and adds an effect size line when "Glass’s delta" is on.
3. Step 1, "Choose variables": the step description reads "Pick the number to compare. Add a grouping variable for group comparisons, or a second measurement for paired data."

•

"Outcome": click in the box and choose the number whose mean is tested or compared. It appears as a chip tagged "number" or "count". The line under the box reads "The number whose mean is tested or compared." - "Grouping variable" (optional): choose a category variable; its chip carries the group count instead of a type tag. The line under the box reads "Optional. Splits the outcome into groups to compare. Leave empty for a one-sample or paired test." - "Paired measurement" (optional): choose the second measurement of the same cases. The line under the box reads "Optional. The second measurement of the same cases, for a paired test against the outcome." - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "region is a category; this role needs a number or a count."

4. Step 2, "Set up the test": the step description reads "Choose the test, the direction, and what prints with it." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" and "Detailed" agree on the test setup: "Test" on "Automatic", "Alternative hypothesis" on "Two-sided", "Variance assumption" on "Welch", and the confidence level at 95%. They differ in one place only: "Detailed" switches "Glass’s delta" on. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in three groups:

- "Test setup": "Test", with "Automatic" (the preset), "One-sample", "Independent samples", "Paired", and "Analysis of variance". The line under it reads "Automatic picks the test that matches the variables chosen. Choose a test to fix it." Below it, "Alternative hypothesis": "Two-sided", "Less (one-sided)", or "Greater (one-sided)"; one-sided alternatives apply to the one-sample and two-group procedures only. Below that, "Variance assumption": "Welch" (the preset, unequal variances allowed) or "Student (equal variances)". The choice reaches the analysis of variance too: "Welch" prints the Welch form of the test, and "Student (equal variances)" prints the classic equal-variance form. - "Precision": "Confidence level", the level of the interval printed for the mean difference: 90%, 95%, or 99%. - "Effect sizes": "Glass’s delta", a switch. The line under it reads "Adds an effect size that measures the difference against the comparison group’s spread." The comparison group is the second group.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Test value: 0 · Follow-up comparisons: Automatic · Random seed: 42". Inside:

- "Test value": the value the mean, or the mean paired difference, is tested against; preset to zero. - "Follow-up comparisons": after an analysis of variance, which pairwise group comparisons print: "Automatic" (the preset; Games-Howell under Welch, Tukey-Kramer under the classic form), "None", or "All pairs, Holm-adjusted". - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7.

Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "score by group · Essentials preset", "score against baseline (paired) · Essentials preset", or "score (one sample) · Essentials preset", or it names the first thing still missing: "Choose an outcome variable to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Group statistics" first, then the test table, then "Effect sizes".

Reading the output

"Group statistics" gives each group's n, mean, standard deviation, and standard error: the inputs to everything after. The worked example's test table reads t = 4.08 with p = 0.004 for the Welch two-sample test, a five-point difference on a standard error of about 1.22. The ANOVA example prints the Welch form of the test: the takeaway reads "Mean score differed across the 3 groups, Welch's F(2, 5.8) = 23.55, p = 0.002", the "Welch's one-way ANOVA" table gives F = 23.55 on 2 and 5.8 degrees of freedom, and the "Post hoc comparisons (Games-Howell)" table then compares each pair of groups (A against C differs by 5.5 points, p = 0.002; A against B by 2.0 points, p = 0.031; B against C by 3.5 points, p = 0.013).

Welch's form does not assume that the groups share one variance, which is why its second degrees of freedom is not a whole number. The classic equal-variance F for these scores is 31.0; under the preset it is not the form the module prints. Set "Variance assumption" to "Student (equal variances)" and the classic form is the one that prints, with F = 31.0.

The "Effect sizes" table prints standardized differences (Cohen's d forms, and Glass's delta when requested) with intervals, so a significant result also gets a size. Under a design, the same reading order holds; the titles name the design-based forms.

Downloads and copying

The full set of seven outputs, subscription permitting, including R and SPSS reproductions of the chosen procedure.

When it refuses or warns

- With no outcome selected, the run stops: "Select an outcome variable."
- A categorical outcome stops the run: comparing means needs a continuous outcome, and the message names the variable's level.
- Selecting both a group and a paired variable draws a warning that the paired procedure takes precedence under "auto"; set "Procedure" explicitly to say what you mean.
- One-sided alternatives requested for ANOVA draw the note that they apply to the one-sample and two-group procedures only.

Module 4: Correlation

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

A correlation matrix across a set of variables: Pearson correlations by default, Spearman automatically when either variable of a pair is ordinal (or by choice), Kendall's tau on request, each with a p-value adjusted across the matrix, a confidence interval, and the pair's n. With "controls" variables declared, partial correlations: the association between each pair with the controls held constant. Under a design, correlations are weighted and the significance uses design-consistent degrees of freedom.

The data it needs

Two to fifty variables in the "variables" role (continuous, count, ordinal, or binary), and up to twenty continuous or count "controls" for partial correlations. Worked example, six cases: x reads 2, 4, 5, 6, 8, 9 and y reads 3, 5, 5, 8, 9, 11.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 4: Correlation". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it counts the correlations among the chosen variables (three variables give three correlations, four give six), and, when control variables are chosen, adds the same correlations with the control variables held constant.

3. Step 1, "Choose variables": the step description reads "Pick two or more variables. Add control variables to see the correlations with their effect removed."

- "Variables": click in the box and choose two or more variables. Each chosen variable appears as a chip with a type tag: "number", "category", "count", or "date". The line under the box reads "Two or more numbers, counts, ordered categories, or two-level categories." - "Control variables" (optional): choose the variables whose effect is to be removed. The line under the box reads "Optional. The correlations are printed with the effect of these variables removed." - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "notes is a text variable; this role needs a number or a count or a category or a two-level category."

4.

Step 2, "Set up the correlations": the step description reads "Choose the coefficient, the handling of missing values, and the adjustment for the many tests." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" sets "Coefficient" to "Automatic", "Missing values" to "Listwise deletion", and "Adjustment for many tests" to "Holm". "Detailed" keeps "Automatic" and the 95% confidence level, and changes two settings: "Missing values" to "Pairwise deletion" and "Adjustment for many tests" to "Bonferroni". Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups:

- "Method": "Coefficient", with "Automatic" (the preset), "Pearson", "Spearman", and "Kendall". The line under it reads "Automatic picks Pearson for numbers and ranks for ordered categories. Choose a coefficient to fix it." Below it, "Missing values": "Listwise deletion" (a case with any missing value in the set is dropped from all pairs) or "Pairwise deletion" (each pair uses its own complete cases, so pair counts can differ). The line under it reads "Listwise uses only cases with every variable present. Pairwise uses every case that has the pair." Below that, "Adjustment for many tests": "Holm" (the preset), "None", "Bonferroni", or "Benjamini-Hochberg (BH)". The line under it reads "With many pairs, some small p-values appear by chance. The adjustment keeps the family of tests honest." - "Precision": "Confidence level", the level of the interval printed for each correlation: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. This module has no advanced options of its own, so the section holds only "Random seed", preset to 42, and its summary line reads "Random seed: 42".

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "four variables · Essentials preset", or names the first thing still missing: "Choose at least two variables to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the correlation matrix with, per pair, the coefficient, the adjusted p-value, the interval, and the pair count.

Reading the output

The worked example's Pearson correlation is r = 0.974 (p = 0.001, 95% CI 0.78 to 1.00): a strong positive association. Read the matrix by row: each cell is one pair. With "Missing-data handling" on "Pairwise deletion", check the n column before comparing cells, since a strong-looking pair may rest on fewer cases. With controls declared, the table is the partial matrix: the same pairs with the controls' linear influence removed from both sides, which is the right reading when a third variable plausibly drives both.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Fewer than two variables stops the run: "Select at least two variables."
- A text variable in the set stops the run and names its level: correlation needs continuous or ordinal variables.
- A pair with no variance (a constant variable) cannot correlate; the checks list flags it rather than printing a number.

Module 5: Linear regression

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

Multiple linear regression with a choice of standard-error estimators (heteroskedasticity-consistent HC3 by default, with classical errors shown alongside), an ANOVA table for the model (Type III by default, the SPSS form), standardized coefficients, collinearity diagnostics, optional two-way interactions, optional hierarchical blocks, estimated marginal means for a categorical predictor, and fitted values with intervals on request. Under a design, estimation is weighted with design-based standard errors and Wald tests.

The data it needs

One continuous "outcome", one to thirty "predictors" of any measured level (categorical predictors expand into indicator columns), and an optional "cluster" variable for cluster-robust errors. Worked example, eight cases: x runs 1 to 8 and y reads 3, 5, 6, 8, 10, 11, 13, 15.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 5: Linear regression". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: it counts the regression, adds coefficients with intervals, fit statistics, and checks, counts any interaction terms in the model, and adds fitted values with intervals, or adjusted means, when those options are on.

3. Step 1, "Choose variables": the step description reads "Pick the outcome and the predictors that explain it. Add interaction pairs here, or a cluster variable when cases sit inside groups."

- "Outcome": click in the box and choose the number the model explains. It appears as a chip tagged "number" or "count". The line under the box reads "The number the model explains." - "Predictors": choose one or more predictors; each appears as a chip with a type tag. The line under the box reads "The variables the outcome is regressed on. Categories enter as indicator variables." - "Cluster variable" (optional): choose the variable that names the group each case belongs to. The line under the box reads "Optional. Names the group each case belongs to, for cluster-robust standard errors." -

"Interactions", in its own group in this step: a chip input with two variable lists and an "Add pair" button. Choose the two predictors of a pair and press "Add pair"; the pair appears as a removable chip, printed with a multiplication sign between the two names. The line under it reads "Add a pair of predictors to include their interaction: whether one predictor’s effect depends on the other." The two main effects are added automatically when missing, and the run prints a check naming any it added. A pair must name two different variables, and a pair that is already in the model is refused. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "notes is a text variable; this role needs a number or a category."

4. Step 2, "Build the model": the step description reads "Choose how the estimates and their precision are computed." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" and "Detailed" agree on the model: no interactions, "Standard errors" on "HC3", "Standardized coefficients" on, no adjusted means factor, the confidence level at 95%, and no blocks. They differ in one place only: "Detailed" switches "Fitted values and intervals" on. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups:

- "Model": "Standard errors", with "Classical", "Robust (HC0)", "Robust (HC1)", "Robust (HC2)", "Robust (HC3)" (the preset), and "Cluster-robust". The line under it reads "Robust standard errors stay valid when the spread of the errors is uneven. HC3 is the recommended robust form; HC0 to HC2 are earlier variants. Cluster-robust needs a cluster variable." Choosing "Cluster-robust" with no cluster variable stops the run: "Cluster-robust standard errors need a cluster variable." Below it, two switches: "Standardized coefficients", which also prints each coefficient in standard-deviation units, so predictors compare on one scale; and "Fitted values and intervals", which adds a table of each case’s fitted value with its confidence and prediction intervals, for checking the fit. Below them, "Adjusted means for a factor": a list drawn from the category predictors you have chosen, or "(none)". The line under it reads "Prints the model’s adjusted mean outcome at each level of one category predictor." - "Precision": "Confidence level", the level of the intervals printed for the coefficients: 90%, 95%, or 99%.

6. Step 3, "Add blocks (optional)": the step description reads "Enter predictors in blocks to watch the fit change as each block joins the model." The one group, "Blocks", is a chip input: type the predictor names in the block, separated by commas, and press "Add block". Each block appears as a removable chip, labeled "Block 1", "Block 2", and so on, with its predictor names after the label. The line under it reads "Add one block at a time: type the predictor names in the block, separated by commas. Blocks enter in order, and the change in fit prints at each step."

7. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Category coding: Dummy (reference) coding · Center number predictors: Off · Analysis-of-variance type: Type III · Random seed: 42". Inside:

- "Category coding": "Dummy (reference) coding" (the preset) or "Deviation coding". The line under it reads "How category predictors are coded for the coefficients. Deviation coding compares each level with the overall mean. The Type III analysis-of-variance table always uses sum-to-zero contrasts, whichever coding is chosen here." - "Center number predictors": a switch. The line under it reads "Subtracts each number predictor’s mean before fitting, which makes the intercept the expected outcome at average values." - "Analysis-of-variance type": "Type III" (the preset) or "Type II". The line under it reads "How the model’s analysis-of-variance table divides the explained variation among predictors." - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

8. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "income on three predictors · Essentials preset", or names the first thing still missing: "Choose an outcome and at least one predictor to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the coefficient table ("Linear regression: <outcome>"), "Model fit", the analysis-of-variance table, and "Collinearity".

Reading the output

The worked example prints an intercept of 1.32 and a slope of 1.68 (SE 0.05), with R squared of 0.996: nearly two units of y per unit of x, tightly estimated. Read the coefficient table for the effects (unstandardized for units, standardized for relative size), "Model fit" for R squared and the residual scale, the ANOVA table for each term's joint test, and "Collinearity" for variance inflation factors: a large factor on a predictor means its coefficient is sharing information with the others and its error bar is honestly wide. Under a design the fit table notes that a weighted R squared is descriptive only.

Downloads and copying

The full set of seven outputs, subscription permitting, including the R script and SPSS syntax for the same specification.

When it refuses or warns

- No outcome or no predictor stops the run with the missing role named.
- "Standard-error type" on "cluster" with no cluster variable stops the run: "Cluster-robust standard errors need a cluster variable."
- An interaction entry that is not a pair ("a:b" with exactly two distinct variables) stops the run and names the entry.

- Perfect collinearity (a predictor that is an exact combination of others) is flagged in the checks and the term is dropped rather than estimated falsely.

Module 6: Logistic regression

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

Binary logistic regression: coefficients on the log-odds scale with odds ratios, Wald or profile-likelihood intervals, a likelihood-ratio test per term, model fit, a classification table at a chosen threshold, and average marginal effects on the probability scale. Separation (a predictor that perfectly splits the outcomes) is handled by a Firth penalized fit, automatically by default, and the output says when it fired. Under a design, coefficients carry design-based standard errors and Wald tests.

The data it needs

One binary "outcome" and one to thirty "predictors". The outcome's event (the level modeled as the one that happens) is the last observed level unless you set it. Worked example, forty cases with a binary treatment predictor: the control group has 6 events in 20 cases, the treated group 12 in 20.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".
2. Choose "Module 6: Logistic regression". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it counts the logistic model and adds odds ratios with intervals and a classification table.
3. Step 1, "Choose variables": the step description reads "Pick the two-level outcome and the predictors that explain it."

- "Outcome": click in the box and choose a two-level category variable. It appears as a chip tagged "category". The line under the box reads "A two-level category. The model gives the probability of the event level." - "Predictors": choose one or more predictors; each appears as a chip with a type tag. The line under the box reads "The variables the event probability is regressed on." - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "score is a number; this role needs a two-level category."

4. Step 2, "Set up the model": the step description reads "Choose how the model is fitted and how its results print." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" and "Detailed" agree on the fitting:

no event level chosen, "Penalized fitting" on "Automatic", the classification cutoff at 0.5, and the confidence level at 95%. They differ in one place only: "Detailed" sets "Interval method" to "Profile likelihood", where "Essentials" leaves it on "Wald". Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups:

- "Model fitting": "Event level", a list drawn from the levels of the chosen outcome. The line under it reads "Which level of the outcome counts as the event whose probability is modeled." Leaving it on "(none)" models the last observed level as the event, and the run warns until you confirm it, because odds ratios read backwards when the event is flipped. Below it, "Penalized fitting": "Automatic" (the preset), "Always", or "Never". The line under it reads "A penalty that keeps estimates steady when one outcome level is rare or a predictor separates the outcome. Automatic applies it when the data call for it." - "Intervals and classification": "Interval method", "Wald" (the preset) or "Profile likelihood". The line under it reads "Profile intervals follow the likelihood and are more accurate with small samples or rare events." Below it, "Classification cutoff": cases with a predicted probability at or above the cutoff are classified as the event; the preset is 0.5. Below that, "Confidence level", the level of the intervals printed for the coefficients and odds ratios: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. This module has no advanced options of its own, so the section holds only "Random seed", preset to 42, and its summary line reads "Random seed: 42".

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "bought on three predictors · Essentials preset", or names the first thing still missing: "Choose an outcome and at least one predictor to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the coefficient table ("Logistic regression: <outcome>"), "Model fit", "Average marginal effects", and the classification table. The coefficient table title names the event in force, in the form "Logistic regression: <outcome> = '<event>'".

Reading the output

The worked example prints an intercept of −0.85 and a treatment coefficient of 1.25 (SE 0.67), an odds ratio of 3.5: the treated group's odds of the event are three and a half times the control group's. The z statistic is 1.88 (p = 0.06), so at this size the estimate is large but not tightly pinned. Read the marginal effects table for the same effect on the probability scale (percentage points), the classification table for how the model sorts the actual cases at the threshold, and the likelihood-ratio table for each term's contribution. When Firth fires, the header notes it and the intervals are the penalized-likelihood kind.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- An outcome without exactly two observed levels stops the run; for more than two levels, Module 14's multinomial model is the tool.
- The module models the last observed outcome level as the event unless you set "Event level (the outcome level modelled as the event)" yourself; the methods paragraph states the event in force ("The event modelled is outcome = \"1\""). Always check it names the event you meant; a flipped event flips every sign in the table.
- Complete separation without Firth (option on "never") fails rather than printing infinite coefficients.

Module 7: Sample size and power planner

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

This module runs no analysis on loaded data; it plans studies. It needs no dataset, though you may load one to borrow pilot values.

What it computes

For thirteen study designs, the third quantity from any two of sample size, power, and detectable effect: one-sample, paired, and two-sample t tests; one proportion against a value; two proportions; one-way ANOVA; correlation; multiple regression (R squared); chi-square; logistic regression; survey sample size for a target margin of error; stratified allocation of a total sample across strata; and the projected relative standard error of a subgroup estimate. It prints a "Plan" table and a sensitivity table of required n at other power levels.

The data it needs

Numbers, not a dataset: the effect you want to be able to detect (or the difference and standard deviation behind it), the significance level, and the target power, plus the design's own inputs (number of groups for ANOVA, degrees of freedom for chi-square, event rate for logistic, strata sizes and standard deviations for allocation, and so on). Pilot variables from loaded data can supply a standard deviation or a proportion.

The steps

1. This planner needs no dataset, and the action bar never asks for one. Load data only if you want to borrow pilot values from it, as Step 1 describes. The "Survey design" card at the side states that the planner does not use a survey design; it works from the design quantities, not from data.
2. Choose "Module 7: Sample size and power planner". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it names the plan and what the run solves for, in the form "plan for a two-sample t test, solving for the sample size".
3. Step 1, "Pilot data (optional)": the step description reads "The planner needs no data. Load pilot data only to borrow its values, such as a standard deviation, for the design quantities." The one picker is "Pilot variables": click in the box and choose variables from the loaded data; each appears as a chip with a type tag. The line under the box reads "Optional. Pick pilot variables to borrow their values for the quantities below." Leave the box empty when you have the quantities from elsewhere.
4.

Step 2, "Design the study": the step description reads "Pick the planned analysis, say what to solve for, and fill in the design quantities." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". Both presets plan a "Two-sample t test", solve for "Sample size", and set "Alternative hypothesis" to "Two-sided" and "Significance level" to 0.05. They differ in one place only: "Detailed" raises "Target power" from 0.8 to 0.9. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The "Design" group holds five controls:

- "Planned analysis": the analysis the study is being planned for. Choose one of thirteen, by its printed name: "One-sample t test", "Paired t test", "Two-sample t test", "Two proportions", "One proportion against a value", "One-way ANOVA", "Correlation", "Multiple regression (R squared)", "Chi-square test", "Logistic regression", "Survey sample size for a margin of error", "Stratified allocation", and "Subgroup projected RSE". The quantities in the next group change with the choice. - "Solve for": a segmented control with "Sample size", "Power", and "Detectable effect". The line under it reads "What the planner works out from the quantities given." - "Alternative hypothesis": "Two-sided" (the preset), "Less (one-sided)", or "Greater (one-sided)", where the design uses one. The line under it reads "Two-sided plans for a difference in either direction. One-sided plans look in one direction only." - "Significance level": the false-positive rate the plan accepts, as a proportion; the usual value is 0.05. - "Target power": the chance of detecting the effect if it is real, as a proportion; the usual target is 0.8. It is used when solving for sample size or for the detectable effect.

6. The "Design quantities" group holds the fields for the planned analysis chosen, and the fields change when that choice changes:

- For a t test: "Effect size (Cohen d)", or a "Difference" with its "Standard deviation"; then "n", or "n per group" for the two-sample test, used when solving for power or for the detectable effect. The two-sample test also takes "Allocation ratio (second group / first group)", preset to one for equal group sizes. - For "Two proportions": "p1 (first proportion)", "p2 (second proportion)", a switch, "Use the arcsine transformation", the allocation ratio, and "n per group". For "One proportion against a value": "Null proportion p0", "Alternative proportion p1", a switch, "Exact binomial calculation", and "n". - For "One-way ANOVA": "Effect size (Cohen f)" and "Number of groups", with "n per group". For "Correlation": "Correlation r" and "n". For "Chi-square test": "Effect size (Cohen w)", "Degrees of freedom", and "Total N". - For "Multiple regression (R squared)": "f squared", "R squared", "Number of predictors", "Predictors added (for an R squared change)", "Total predictors (with the added set)", and "Total N". - For "Logistic regression": "Odds ratio", "Overall event rate", "Predictor type" ("Normal" for a continuous predictor, or "Binary" for a two-level predictor, with its prevalence below), "Predictor prevalence (share with the predictor)", "R squared for the other predictors", and "n". - For "Survey sample size for a margin of error": "Anticipated proportion",

"Target margin of error" (as a proportion: 0.03 for plus or minus three points), "Population size (optional)" (when given, the finite population correction applies), "Design effect" (preset to one; or give an "Intraclass correlation" and an "Average cluster size" and the design effect is computed from them), "Anticipated response rate" (as a proportion: 0.7 for seventy percent; the fielded size is inflated to still land the completes), and "n", used when solving for the achieved margin of error. - For "Stratified allocation": "Stratum population sizes, comma-separated", "Stratum standard deviations, comma-separated" (one per stratum, in the same order), "Stratum costs, comma-separated (optional)" (one per stratum; the cost-optimal allocation uses them), and "Total sample size to allocate". - For "Subgroup projected RSE": "Subgroup proportion", "Subgroup sample size", and "Design effect", preset to one.

7. "Advanced settings" stays closed unless you need it. This module has no advanced options of its own, so the section holds only "Random seed", preset to 42, and its summary line reads "Random seed: 42".

8. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as "study plan · Essentials preset" (with the preset name you are on), or names the first thing still missing: "Fill in the design quantities to run". The run is blocked until at least one design quantity is given; after that, a design whose inputs are still incomplete stops under "Check the setup" and names the missing quantity. "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the "Plan" table with the solved quantity, then the sensitivity table.

Reading the output

Worked example: "Two-sample t test", solving for n, effect size d of 0.5, alpha 0.05, power 0.8, two-sided. The plan calls for 64 per group, 128 in total (the exact solution is 63.77 per group before rounding up). Read the sensitivity table as the price of certainty: it prices the same study at other power levels, which is the fastest way to negotiate a design with a budget. For the survey procedure, the plan accounts for the design effect and response rate you declare; for allocation, the "Stratum allocations" table divides the total by the Neyman rule (more sample where a stratum is larger or more variable, tempered by cost when costs are given).

A limit to know (the module prints it where it applies): the logistic procedure's formula carries a calibrated accuracy note from the module's own validation, because closed-form logistic power is approximate at small effects. Treat its answer as a planning figure, and size up when the study is expensive to repeat.

Downloads and copying

The full set of seven outputs, subscription permitting. The R script reproduces the calculation in the pwr family of functions, which makes the plan easy to hand to a collaborator.

When it refuses or warns

- A design whose inputs are incomplete stops under "Check the setup" and names the missing quantity, for example "Provide an effect size (Cohen d) or a difference with a standard deviation.", "Provide p1 in (0, 1).", or "Provide the total sample size to allocate."
- Values outside their ranges (a proportion at 0 or 1, an odds ratio of exactly 1, a correlation of 0) are refused with the range stated, because the solved quantity would be infinite or undefined there.

Module 8: Rank-Based Tests and Exact Alternatives

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

The nonparametric counterparts of the mean comparisons, plus exact alternatives: Mann-Whitney (two independent groups), Wilcoxon signed-rank (paired), the sign test, Kruskal-Wallis (several groups) with Dunn follow-up comparisons, Friedman (repeated measures), McNemar (paired binary), Cochran's Q (repeated binary), and a design-based rank test when a survey design is in force. Hodges-Lehmann location estimates with intervals accompany the two-sample and paired procedures.

The data it needs

One "outcome" plus a "group" variable (independent designs) or a "paired" variable (paired designs), or three or more "repeated" measurement variables (Friedman, Cochran). Worked example: group A reads 12, 15, 18, 20 and group B reads 9, 11, 14, 16, 19.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 8: Rank-Based Tests and Exact Alternatives". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once an outcome or a set of repeated measures is in place, it lists one rank-based test with its effect size, and the group medians and Hodges-Lehmann intervals where the procedure gives them.

3. Step 1, "Choose variables": the step description reads "The layout you pick decides the test: an outcome with a grouping variable, an outcome with a paired variable, or a set of repeated measures."

- "Outcome": click in the box and choose the measure being compared. It appears as a chip tagged with its type. The line under the box reads "The measure being compared. Leave empty when you pick repeated measures instead." - "Grouping variable" (optional): choose a category variable; its chip carries the group count instead of a type tag. The line under the box reads "Optional. Splits the outcome into groups: two groups give a Mann-Whitney test, more give a Kruskal-Wallis test." - "Paired variable" (optional): choose the second measurement of a pair, matched row by row with the outcome. - "Repeated measures" (optional): choose three or more measurements of the

same cases, for a Friedman test. When these are picked, the outcome stays empty. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "region is a category; this role needs a number or a count."

4. Step 2, "Set up the test": the step description reads "Leave the procedure on Automatic to let the layout of the data choose it, or name the test yourself." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" and "Detailed" agree on everything except the confidence level: 95% in "Essentials", 99% in "Detailed". Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups:

- "Procedure": "Procedure", with "Automatic" (the preset), "Mann-Whitney U", "Wilcoxon signed-rank", "Sign test", "Kruskal-Wallis", "Friedman", "McNemar", "Cochran's Q", and "Design-based rank test". The line under it reads "Automatic chooses from the layout of the data, the design, and whether the measures are two-level. Naming a procedure runs exactly that test." Below it, "Test value": the value the one-sample procedures test the outcome against; preset to zero. - "Precision": "Confidence level", the level of the Hodges-Lehmann intervals printed with the location tests: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. It holds only the "Random seed", preset to 42, and its summary line states the current value, for example "Random seed: 42".

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "score by region · Essentials preset", or names the first thing still missing: "Choose an outcome or repeated measures to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Group rank summaries" with mean ranks, then the test table "Rank-based tests".

Reading the output

The worked example prints U = 14 with a two-sided p of about 0.41: on these small samples the rank evidence does not separate the groups, whatever the raw gap suggests. Rank tests compare distributions through ranks, so read the mean ranks beside the statistic. After a Kruskal-Wallis result, the "Dunn follow-up comparisons (Holm adjusted)" table tests each pair with the adjustment applied. McNemar prints the paired 2 by 2 table itself, so the discordant pairs that drive the test are visible.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Selecting both a grouping and a paired variable stops the run: choose one structure.
- "Repeated-measures procedures need three or more measurement variables.", and they stand alone: remove grouping and paired variables.
- McNemar with non-binary variables stops the run: "McNemar needs two binary variables."
- A named procedure whose structure is missing stops the run naming what it needs (an outcome and a grouping variable, or paired variables, as the case may be).

Module 9: Key driver analysis

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

A decomposition of a regression's R squared across the predictors ("drivers"), so each driver gets a share of the explained variation that sums to the whole: the LMG method (the Shapley decomposition, averaging each driver's contribution over every ordering) or Johnson's relative weights (the automatic fallback when the driver count makes the full enumeration infeasible). Bootstrap intervals accompany the shares, and each driver also gets a performance score, so importance can be read against performance in the classic quadrant arrangement.

The data it needs

One "outcome" (continuous, ordinal with five or more points treated as continuous, or binary) and two to thirty numeric "predictors". Worked example, twelve respondents rating an outcome y and three drivers; the full data are: d1 = 3, 5, 4, 6, 5, 7, 4, 6, 5, 8, 6, 7; d2 = 2, 3, 3, 4, 3, 5, 4, 4, 5, 5, 4, 6; d3 = 6, 5, 7, 4, 6, 5, 7, 6, 4, 5, 6, 5; y = 10.3, 14.1, 13.4, 17.4, 14.3, 20.9, 13.9, 18.1, 15.9, 23.0, 17.4, 21.7.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 9: Key driver analysis". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once an outcome and at least two drivers are in place, it counts the drivers ranked by importance, with intervals, and adds the performance for each driver and the quadrant chart.

3. Step 1, "Choose variables": the step description reads "Pick the outcome the drivers should explain, and the drivers themselves."

- "Outcome": click in the box and choose the measure the drivers are judged against, such as overall satisfaction. It appears as a chip tagged with its type. - "Drivers": choose the measures whose importance is compared. Pick at least two and at most 30. Each appears as a chip tagged with its type. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4.

Step 2, "Set up the analysis": the step description reads "Choose how importance is divided among the drivers and how each driver's performance is scored." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets agree on the method ("Automatic"), the performance measure ("Mean rescaled to 0-100"), and empty custom splits; they differ in the quadrant split lines: "Essentials" divides at the medians, "Detailed" at the means. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in three groups:

- "Method": "Importance method", with "Automatic" (the preset), "LMG (Shapley)", and "Johnson's relative weights". The line under it reads "Automatic uses the Shapley division with 15 or fewer drivers and relative weights above that. Either method can be named directly." - "Performance": "Performance measure", with "Mean rescaled to 0-100" (the preset) and "Top-box percentage". The line under it reads "How each driver is scored on the quadrant chart: its mean rescaled to 0 to 100, or the share of top-box answers." Below it, "Top-box cut": the smallest value counted as a top-box answer, needed only when performance is the top-box percentage; leave it empty otherwise. - "Quadrant chart": "Quadrant split lines", with "Medians", "Means", and "Custom values". The line under it reads "Where the quadrant chart divides high from low: at the medians, at the means, or at values you give." Below it, "Custom performance split" and "Custom importance split" (the second as a share) hold your own dividing lines, used when the split lines are custom values.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Bootstrap replications: 1,000 · Random seed: 42". Inside:

- "Bootstrap replications": how many bootstrap samples draw the intervals around the importance shares, from 500 to 5,000; preset to 1,000. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "4 drivers of overall · Essentials preset", or names the first thing still missing: "Choose an outcome and at least two drivers to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Key driver analysis" titled with your outcome, with the drivers ranked by importance and the quadrant chart.

Reading the output

The worked example's model explains 99.1% of the variation, and the importance shares come out d1 at 56.8%, d2 at 36.1%, and d3 at 7.1% of the explained variation: the first driver does most of the work, and improving the third would buy little. Read shares against their bootstrap intervals before ranking two close drivers, and read importance against

performance in the same table: high importance with low performance is where effort pays; high importance with high performance is what to protect. The context table (plain correlations and standardized coefficients) is printed so a share that surprises you can be traced. Under a declared design the regression behind the shares is the design-based fit, and the output notes it.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Fewer than two drivers, a missing outcome, or the outcome also listed as a driver stops the run with the reason named.
- Top-box performance without a cut stops the run: "Top-box performance requires the top-box cut."
- A bootstrap count outside 500 to 5,000 stops the run; an LMG request on too many drivers is refused with the count printed and Johnson named as the route.
- An ordinal outcome with five or more points draws the printed note that it is treated as continuous.

Module 10: A/B test analyzer

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

A complete readout of a randomized experiment: arm sizes and rates, each variant against the control (or all pairs) with multiplicity adjustment, binary metrics by difference with score intervals plus a Bayesian probability-to-beat summary, continuous metrics by Welch's t (plain, winsorized, or bootstrap), count metrics, ratio metrics from numerator and denominator variables, an optional CUPED variance reduction using a pre-experiment covariate, a sample-ratio-mismatch (SRM) check against the planned allocation, an early-look warning when the observed n falls short of the planned n, and cumulative-over-time views when a timestamp is supplied. This module analyzes an experiment you ran; it declares no survey design.

The data it needs

One "variant" variable (the assignment), one primary "metric", and optionally "secondaryMetrics" (up to twelve), one "covariate" (pre-experiment, for CUPED), "numerator" and "denominator" for a ratio metric, and a "timestamp". Worked example: arm A (control) converts 24 of 200 visitors (12.0%), arm B converts 36 of 200 (18.0%).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 10: A/B test analyzer". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: once a variant and a metric are in place, it lists one comparison table for the primary metric, a line for any secondary metrics, and a decision summary in plain words.

3. Step 1, "Choose variables": the step description reads "Pick the variant that splits the cases and the metric that decides the result."

- "Variant": click in the box and choose the experiment arm each case belongs to. Its chip carries the group count instead of a type tag. - "Primary metric": choose the measure the experiment is judged on. It appears as a chip tagged with its type. - "Secondary metrics" (optional): further measures, each compared the same way as the primary. Up to 12. - "Pre-experiment covariate" (optional): a measure taken before the experiment, used by the CUPED variance reduction. - "Numerator" (optional) and "Denominator"

(optional): the top and bottom halves of a ratio metric, such as revenue over sessions. - "Timestamp" (optional): when each case arrived, for the early-look warning. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the analysis": the step description reads "Say what kind of metric the primary is and which comparisons to print." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets agree on the whole setup except two places: "Detailed" compares all pairs (where "Essentials" compares each variant against the control) and uses the bootstrap interval for continuous metrics (where "Essentials" uses Welch). Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The Step 2 options sit in two groups:

- "Primary metric": "Primary metric type", with "Automatic" (the preset), "Binary", "Continuous", "Count", and "Ratio". The line under it reads "Automatic treats a 0 and 1 metric as two-level and anything else as a continuous measure. Ratio uses the numerator and denominator variables." Below it, "Control variant": which variant level is the control the others are compared against; the first observed level is the preset. - "Comparisons": "Comparisons", with "Each variant against control" (the preset) and "All pairs". Below it, "Adjustment for many comparisons", with "Holm" (the preset), "Benjamini-Hochberg (BH)", and "None". The line under it reads "How the p-values within each metric family are adjusted when several comparisons print."

6. Step 3, "Methods for the metrics": the step description reads "Choose the interval and variance methods the comparisons use." The options sit in three groups:

- "Intervals": "Interval for a difference in shares", with "Newcombe score" (the preset) and "Wald". The line under it reads "The interval method for two-level metrics: the Newcombe score interval or the simpler Wald interval." Next, "Method for continuous metrics", with "Welch" (the preset), "Winsorized Welch", and "Bootstrap BCa". Next, "Confidence level": 90%, 95%, or 99%. - "Variance reduction": "Variance reduction", with "None" (the preset) and "CUPED". The line under it reads "CUPED uses the pre-experiment covariate to tighten the comparison for continuous and count primaries." - One switch: "Bayesian summary". The line under it reads "For two-level metrics, add the chance that each variant beats the reference arm and the expected loss of choosing it." It starts on.

7. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Winsorizing percentile: 99 · Prior alpha: 1 · Prior beta: 1 · Minimum effect of interest: Not set · Planned allocation shares: Not set · Planned total sample size: Not set · Analysis unit: Randomization unit · Random seed: 42". Inside:

-

"Winsorizing percentile": the percentile at which the winsorized Welch method caps the metric values; preset to 99. - "Prior alpha" and "Prior beta": the two halves of the beta prior behind the Bayesian summary; each preset to 1. - "Minimum effect of interest": the smallest difference that would matter, used only in the decision summary. Leave empty to skip it. - "Planned allocation shares": the planned share of each variant, in variant order, separated by commas, such as 0.5, 0.5. Used by the sample-ratio check. Leave empty to skip it. - "Planned total sample size": the planned number of cases. A smaller observed count raises the early-look warning. Leave empty to skip it. - "Analysis unit": "Randomization unit" (the preset) or "Other". The line under it reads "The unit the randomization assigned. Anything else needs a ratio metric to define it." - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

8. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "converted by arm · Essentials preset", or names the first thing still missing: "Choose a variant and a metric to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the comparison table for the primary metric, then the decision summary.

Reading the output

The worked example's primary difference is 6.0 percentage points with a Newcombe interval of −1.0 to 13.0 points and an unadjusted p of about 0.09; the Bayesian summary puts the probability that B beats A at about 0.95. Read the adjusted p-values when several variants or metrics are in play, and read the checks first: an SRM flag means the realized allocation departs from the plan beyond chance, which usually means the randomization or the logging needs attention before the effect estimate is trusted; an early-look warning means the run stopped short of the planned n, so a "not significant" read is a statement about power, not about the effect.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- A missing variant, a variant with fewer than two levels, or a control level that is not an observed level stops the run, naming the problem.
- "CUPED requires the pre-experiment covariate variable." and "A ratio metric requires numerator and denominator variables." are hard stops for those options.
- Planned allocation shares that are not positive, or that do not match the variant's level count, stop the run.
- Priors must be positive, and the winsorizing percentile must sit between 50 and 100.

Module 11: Survey weighting and raking

Access: free to run on screen at all times, for the weighting procedure. Downloads and "Copy Results" need the subscription. The second procedure in this chapter (propensity adjustment) follows the subscription access line given in its own section.

What it computes

A new weight variable for your dataset, built in stages: base weights (your declared design weights, or a variable you supply), an optional nonresponse adjustment (by weighting classes, or by a propensity model), and calibration to population targets: raking across separate margins (the preset), post-stratification to one joint classification, or generalized regression calibration (linear, or bounded logit with factors kept inside stated bounds). It prints diagnostics for the weights it built, checks each margin's achieved totals against its targets, and compares key variables' means unweighted, base-weighted, and final-weighted, so the effect of the weighting is visible before you use it.

The data it needs

Your dataset, base weights (declared with the design or named directly), and one margin per calibration variable with a target for every observed category. Worked example, forty respondents in three columns. `sex` is M for the first twenty-two rows and F for the last eighteen. `bw`, the base weight, is 1 for every row. `key`, a satisfaction score, for the twenty-two men: 42.5, 57.5, 44.5, 55.5, 46.5, 53.5, 48.5, 51.5, 49.5, 50.5, 45.5, 54.5, 47.5, 52.5, 43.5, 56.5, 49.0, 51.0, 48.0, 52.0, 50.0, 50.0 (averaging 50). For the eighteen women: 52.5, 67.5, 54.5, 65.5, 56.5, 63.5, 58.5, 61.5, 59.5, 60.5, 55.5, 64.5, 57.5, 62.5, 59.0, 61.0, 60.0, 60.0 (averaging 60). The population targets are 480 men and 520 women in a population of 1,000.

The steps

1. Load data. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to the "Survey design" card; that card matters here, because the base weights can start from the weight declared there. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 11: Survey weighting and raking". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: it lists a new version of your data with the calibrated weight added, counts the calibration margins matched once margins exist, and adds a weighting report with the weight distribution and diagnostics.

3. Step 1, "Choose variables": the step description reads "Key variables are optional. The report compares each key variable’s mean before and after weighting, so you can see what the calibration changed." The one picker is "Key variables (optional)": click in the

box and choose one or more number or count variables; each appears as a chip tagged "number" or "count". The line under the box reads "Optional. For each key variable, the report compares its unweighted mean with its base-weighted and final-weighted means. Key variables play no part in computing the weights." Leave the box empty to run without any.

4. Step 2, "Set up the weighting": the step description reads "Choose the starting weights and how they are adjusted to match known totals." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" sets "Calibration method" to "Raking", "Proportions scaled to" to "Population", leaves "Match the first margin’s total" off, and names the new weight "weight_final". "Detailed" changes two of those: "Calibration method" to "GREG", and "Match the first margin’s total" on. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The "Weighting setup" group holds four fields:

- "Base weights": choose "Declared design weights" to start from the weight declared in "Survey design", or choose the variable holding the base weights. The line under it reads "The starting weights the calibration adjusts: the design weight declared under Survey design, or a weight variable already in the data." The worked example in "The data it needs" picks the base weight variable given there. - "Calibration method": "Raking" (the preset), "Post-stratification", "GREG", or "GREG (logit)". The line under it reads "How the weights are adjusted so the weighted data matches the targets." Under "GREG (logit)", the calibration factors are kept between 0.5 and 2 by default. - "Proportions scaled to": "Population" (the preset) or "Sample". The line under it reads "Only margins entered as proportions use this: their targets are scaled to the population total (the base-weighted total) or to the respondent count. Margins entered as totals are used exactly as entered." - "Match the first margin’s total": a switch. The line under it reads "Rescale every later margin so its targets sum to the same total as the first margin." Turn it on when your margins come from sources with differing totals and the first is the one to trust.

6. Below that group, the "New weight" group holds one field, "Name for the new weight": the name the calibrated weight is written to a new version of your data under. The field is pre-filled with "weight_final"; type another name to use it instead. The line under it reads "The calibrated weight is written to a new version of your data under this name; the loaded data are not modified."

7. Step 3, "Add calibration margins": the step description reads "For each margin variable, enter the known total or share for every category. The weights are adjusted until the weighted data matches." The one group, "Calibration margins", is a builder:

-

"Margin variable": choose a category variable whose distribution the weights should match. - "Target scale": "Total" (the preset) or "Proportion". The line under it reads "Enter this margin’s targets as population totals or as proportions of the total." - One target input then appears for each category of the chosen variable, labeled "<category> target", with that category's respondent count beside it. Type a target into every category field: a margin missing a target for an observed category cannot run, and proportion targets must sum to one. - Press "Add Margin" to add the margin. Each added margin appears as a removable chip in the list under the button. Repeat the builder for each further variable. A margin for a variable can be edited only by removing it and adding it again.

8. "Advanced settings" stays closed unless you need it. This module has no advanced options of its own, so the section holds only "Random seed", preset to 42, and its summary line reads "Random seed: 42".

9. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "weighting · Essentials preset", or "means compared for <key variable> · Essentials preset" when key variables are chosen, or it names the first thing still missing: "Add at least one complete margin to run". The run stays blocked until at least one complete margin exists. "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Weight diagnostics", one "Margin check" table per margin, the "Key variables" comparison when key variables are chosen, and the new weight variable's table.

Reading the output

The worked example (a single margin, so raking and post-stratification agree) gives every man a final weight of 480/22 = 21.82 and every woman 520/18 = 28.89; the final weights sum to 1,000. The key-variable table shows the outcome's mean moving from 54.5 unweighted to 55.2 final-weighted: the weighting pulls the estimate toward the women's higher average, because women were under-sampled relative to the population. "Weight diagnostics" reports the weight distribution (minimum, maximum, and the variation that costs precision); the margin checks should show achieved totals equal to targets, and any margin that could not be hit exactly says so rather than passing silently.

The new weight itself appears as a table in the results, one row per respondent ("New weight variable "weight_final" by row"). Below the results, a "New dataset" block delivers the merged data itself: a line naming the new weight column and the row count, ending "The original data are unchanged.", a preview of the first rows (your first two variables and the new weight column), and two buttons. "Download the new dataset (CSV)" saves the merged file; like every download, it waits for a verified license key. "Use this dataset now" loads the merged data as the current dataset, so the new weight can be declared in "Survey design" and used by later modules in this session.

Downloads and copying

The full set of seven outputs, subscription permitting, including the new weight variable by row. The merged dataset itself travels through the "New dataset" block below the results, described under "Reading the output".

When it refuses or warns

- No base weights specified, no margin declared, or a margin missing a target for an observed category stops the run and names the gap: "Targets must cover every observed category (recode first if a category should be collapsed)."
- Proportion targets that do not sum to 1 stop the run, with the actual sum printed.
- Post-stratification with more than one margin stops the run: it calibrates to one joint classification; cross the variables into one joint variable first, or use raking.
- Bounded logit calibration with bounds that do not satisfy 0 < lower < 1 < upper stops the run and prints the bounds it received.
- The nonresponse adjustments refuse without their inputs: a class variable for the weighting-class form, at least one predictor for the propensity form, and the response variable must exist.

Module 11, second procedure: Propensity adjustment for nonprobability samples

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this procedure needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Pseudo-weights for a nonprobability sample (an opt-in panel, a web convenience sample) anchored to a reference probability sample. The two samples are stacked, a model estimates each case's propensity (its modeled chance of being in the nonprobability sample given shared covariates), and each nonprobability case receives a pseudo-weight from that propensity, in the odds form $(1 - p)/p$ by default or the inverse form $1/p$. Optionally the pseudo-weights are calibrated to the reference sample's margins afterward. Diagnostics and a covariate balance table (nonprobability against reference, design-weighted) show whether the adjustment brought the samples into line.

The data it needs

Two datasets in one working session: the nonprobability sample loaded as your data, with the reference probability sample and its declared design available to the procedure (the run refuses without them), and at least one covariate present, with the same meaning, in both samples. Illustrative arithmetic: a case whose modeled propensity is 0.25 receives an odds-form pseudo-weight of $(1 - 0.25)/0.25 = 3.0$; a case at 0.5 receives 1.0; a case at 0.8 receives 0.25. Cases that look rare in the nonprobability sample count for more, which is the point of the adjustment.

The steps

1. Load the nonprobability sample data first. This procedure uses a reference survey instead of a survey design: load the reference dataset in Step 2 below. The "Survey design" card at the side of the setup explains that this module carries its reference in Step 2. Until data is loaded, the action bar at the bottom reads "Load data to begin".
2. Choose "Module 11, second procedure: Propensity adjustment for nonprobability samples". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: once covariates are picked, it lists one new version of your data with the calibrated weight added, and a report with the weight distribution and diagnostics.
3. Step 1, "Choose variables": the step description reads "Pick the variables both datasets share. The propensity model separates the sample from the reference using these."

•

"Shared covariates": click in the box and choose the variables measured the same way in the sample and the reference. Each appears as a chip tagged with its type. - A variable that cannot serve stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Load the reference sample": the step description reads "The reference probability sample anchors the adjustment: its own design weights say how it represents the population." The fields sit in a "Reference sample" group, with two further groups that appear only for the designs that need them:

- "Reference file": the reference probability sample as a delimited file (CSV or TSV); its covariates carry the same names as in your data. The status line under it reports what is loaded, for example "Reference sample loaded: 80 rows, 6 variables". - "Reference design": "Weights only" (the preset), "Taylor (strata and clusters)", "Replicate weights", or "No design". The line under it reads "The reference sample's declared design. The procedure requires one: its own weights, strata, or clusters." - "Reference weight column": the column holding the reference sample's design weights. Beside it, "Reference weight type": "Sampling", "Frequency", or "Analytic", what the reference weights mean. - With a Taylor design, a "Reference design (Taylor)" group adds the "Reference strata column", the "Reference cluster column", and an optional "Reference population counts column". With replicate weights, a "Reference design (replicate)" group adds the "Reference replicate weight columns", the replicate type ("JK1", "JKn", "BRR", "Fay", "bootstrap", or "other"), the overall scale, the scale for each replicate set, and the Fay coefficient where it applies.

5. Step 3, "Build the pseudo-weights": the step description reads "Choose the predictors and the form of the pseudo-weights, and whether they are calibrated to the reference." The options sit in three groups, with one switch below them:

- "Propensity model": "Predictors", a list holding the covariates chosen in Step 1. The line under it reads "The predictors the propensity model uses, chosen from the covariates selected above. Leave empty to use every covariate." - "Pseudo-weights": "Pseudo-weight form", with "Odds form, (1 − p) / p" (the preset) and "Inverse form, 1 / p". The line under it reads "Odds form gives each case the odds form of its fitted membership chance. Inverse form gives one over that chance." - "Output": "Name for the new weight": the pseudo-weight is written to a new version of your data under this name; the loaded data are not modified. Preset to "pseudo_weight". - "Calibrate to the reference": a switch that starts on. The line under it reads "Adjust the pseudo-weights so the sample matches the reference on the categorical margins and the continuous means."

6. There is no "Advanced settings" fold on this screen beyond the "Random seed", preset to 42, whose current value its summary line states.

7.

Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "4 covariates · Essentials preset", or names the first thing still missing: "Choose covariates to run", or, while the reference is missing, "Load a reference survey and name its weight column to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the report titled "Propensity adjustment for nonprobability samples", with the weight distribution and diagnostics.

Reading the output

Read the balance table first: standardized differences between the nonprobability and reference samples, before and after adjustment, should shrink toward zero on the covariates the adjustment used. The diagnostics table reports the pseudo-weight distribution; a few enormous weights mean a few cases are carrying the sample, which the calibration step tempers but does not repeal. Like Module 11's weights, the pseudo-weight prints as a by-row table in the results, and the same "New dataset" block appears below the results: a preview of the merged data, "Download the new dataset (CSV)" to save it, and "Use this dataset now" to load it as the current dataset for later modules.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Without a reference probability sample carrying a declared design, the run refuses outright: "A reference probability sample with a declared design is required."
- A covariate missing from either sample stops the run and names the sample it is missing from; a predictor that is not among the covariates stops it as well.
- No covariate selected stops the run: "Select at least one covariate shared by both samples."

Module 12: Trend analysis across survey waves

Access: free to run on screen at all times. Downloads and "Copy Results" need the subscription.

What it computes

Change across repeated survey waves: an estimate per wave with its interval, the change between waves, and a formal trend test in the form used for youth risk surveillance reporting (orthogonal polynomial contrasts on the logit scale for a binary outcome: linear and quadratic terms by default, cubic with six or more waves when you ask). Subgroup trends print as parallel blocks. The module refuses to run until you confirm the measure is comparable across waves, and it can stamp every output with a method-change note when wording or weighting changed mid-series.

The data it needs

One "time" variable (the wave), one "outcome", and optionally "covariates" and one "subgroup". Worked example, a binary outcome over three waves: 60 of 200 in the first wave (30.0%), 71 of 210 in the second (33.8%), and 80 of 205 in the third (39.0%).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 12: Trend analysis across survey waves". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: it counts the trend in the outcome across the waves, adds the estimate for each wave, with intervals, and, when a subgroup variable is chosen, adds a separate trend for each subgroup level.

3. Step 1, "Choose variables": the step description reads "Pick the wave variable and the outcome to follow across waves."

- "Wave (time) variable": click in the box and choose the variable that says which wave each case belongs to. It appears as a chip with a type tag. The line under the box reads "The variable that says which wave each case belongs to." - "Outcome": choose the measure whose change across waves is estimated. The line under the box reads "The measure whose change across waves is estimated." - "Adjustment variables" (optional): choose variables the trend is adjusted for. The line under the box reads "Optional. The trend is adjusted for changes in these variables across waves." - "Subgroup variable" (optional): choose a category variable; its chip carries the group count instead of a type tag. The line under the box reads "Optional. A separate trend is printed for each level." -

A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name, for example "notes is a text variable; this role needs a number or a category."

4. Step 2, "Set up the trend": the step description reads "Choose the trend shape and confirm the outcome means the same thing in every wave." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". "Essentials" and "Detailed" agree on the rest of the setup: no wave values added, both comparability switches off, and the confidence level at 95%. They differ in one place only: "Detailed" sets "Trend shape" to "Linear + quadratic + cubic", where "Essentials" leaves it on "Linear + quadratic". Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in three groups:

- "Trend model": "Trend shape", with "Linear + quadratic" (the preset) or "Linear + quadratic + cubic". The line under it reads "Quadratic allows the trend to bend once. Cubic allows two bends." The cubic term is for six or more waves. Below it, "Wave values": a chip input. When the wave variable stores labels instead of numbers, add the numeric value of each wave, one at a time, in wave order; each appears as a removable chip. The line under it reads "Optional. When the wave variable stores labels instead of numbers, add the numeric value of each wave, in wave order." - "Comparability": two switches. "The outcome is comparable across waves" confirms the outcome was asked and measured the same way in every wave; the line under it reads "Confirm the outcome was asked and measured the same way in every wave. The run is blocked until this is confirmed." Below it, "Wording or method changed between waves": the line under it reads "Prints a caution on the output that a change in wording or method may sit behind part of the trend." With the switch on, the note appears on every output. - "Precision": "Confidence level", the level of the intervals printed for the wave estimates: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. This module has no advanced options of its own, so the section holds only "Random seed", preset to 42, and its summary line reads "Random seed: 42".

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "<outcome> across waves · Essentials preset", or names the first thing still missing: "Choose a wave variable and an outcome to run". The comparability switch blocks the run in a second way: without that confirmation, the module refuses to trend a measure you have not vouched for. "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Estimates by wave", "Change between waves", and the trend test tables.

Reading the output

The worked example's by-wave table reads 30.0% (95% CI 24.0 to 36.8), 33.8% (27.7 to 40.5), and 39.0% (32.5 to 45.9). The linear trend term is the headline: a steady climb, rather than a one-wave jump, is what it detects, and the quadratic term tells you whether the climb bends. Read the change-between-waves table for the endpoint story (a nine-point rise across the series). By-subgroup blocks repeat the same tables per subgroup so trends can be compared honestly.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Without the comparability confirmation, the run refuses: the module will not trend a measure you have not vouched for.
- Fewer than two waves, a missing time variable, or a missing outcome stops the run with the gap named.
- With "Question wording or weighting method changed" on, the warning rides on every table rather than once, by design: a series with a break in it should never be quoted without the break mentioned.

Module 13: Scale reliability and exploratory factor analysis

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

For a set of questionnaire items: reliability (Cronbach's alpha with a confidence interval, and the item statistics behind it, including alpha-if-item-deleted), then an exploratory factor analysis: a correlation matrix (Pearson, or polychoric for few-category items), extraction by principal axis, maximum likelihood, or minimum residual, the number of factors chosen by parallel analysis (preset), the Kaiser rule, or a fixed count, a rotation (oblimin preset, promax, or varimax), pattern and structure matrices, factor correlations under oblique rotations, variance explained, and factor scores on request. Weights of the sampling and frequency kinds enter the computations; Taylor and replicate designs are not supported here, and the run says so.

The data it needs

Three to sixty "items", all measured variables. Worked example, six respondents answering four items on a one-to-five scale: respondent rows read (4, 4, 5, 4), (5, 5, 4, 5), (2, 3, 2, 3), (3, 2, 3, 2), (5, 4, 5, 5), (1, 2, 1, 2).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 13: Scale reliability and exploratory factor analysis". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once at least three items are picked, it lists one reliability table for the scale, with the Feldt interval, and the factor loadings and the share of variance each factor carries.

3. Step 1, "Choose items": the step description reads "Pick the items of the scale, at least three."

- "Items": click in the box and choose the scale items. Each appears as a chip tagged with its type. The line under the box reads "The scale items, scored in the same direction. At least three and at most 60." - A variable that cannot serve stays visible in the list, dimmed, with the reason stated under its name, for example "region is a category; items need numbers or ordered answers."

4.

Step 2, "Set up the analysis": the step description reads "Choose the correlation matrix the analysis works from and how the factors are found." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets differ in two places: the correlation matrix (Pearson in "Essentials", polychoric in "Detailed") and the factor scores (off in "Essentials", regression in "Detailed"). Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in four groups:

- "Correlations": "Correlation matrix", with "Pearson" (the preset) and "Polychoric". The line under it reads "Polychoric correlations suit items with five or fewer ordered categories. Pearson suits items scored as numbers." - "Factor analysis": "Extraction method", with "Principal axis" (the preset), "Maximum likelihood", and "Minimum residual". The line under it reads "How the factors are extracted from the correlation matrix." Below it, "Number of factors", with "Parallel analysis" (the preset), "Kaiser criterion (eigenvalue over 1)", and "Fixed number". The line under it reads "Parallel analysis compares the factors against chance. The Kaiser rule keeps factors with an eigenvalue over 1. Fixed uses the count below." Below that, "Factor count": how many factors to extract when the number of factors is fixed; preset to one. Below that, "Rotation", with "Oblimin" (the preset), "Promax", and "Varimax". The line under it reads "How the factor axes are turned to make the loadings readable." Below that, "Factor scores", with "Off" (the preset), "Regression", and "Bartlett". The line under it reads "Whether a score per case is computed for each factor, and by which method." - "Display": "Loading display threshold": loadings smaller than this are left blank in the printed table; preset to 0.3. - "Precision": "Confidence level": the level of the Feldt interval printed for the reliability coefficient: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Items to reverse-code: Q3 · Random seed: 42". Inside:

- "Items to reverse-code": a chip input. Type an item name and press "Add"; each item appears as a removable chip. The line under it reads "Items scored against the scale's direction, reversed before the analysis. Reversal is never automatic: name each item to reverse." - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "reliability of 6 items · Essentials preset", or names the first thing still missing: "Choose at least three items to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Scale reliability and factor analysis" first, then the item statistics and the factor solution.

Reading the output

The worked example's reliability table reads alpha = 0.955, with item means of 3.33, 3.33, 3.33, and 3.50 and a total-score mean of 13.5: the four items hang together tightly. In "Item statistics", an item whose alpha-if-deleted rises above the overall alpha is hurting the scale and is the first candidate for revision. In the factor tables, read the pattern matrix for each item's home factor (loadings above the threshold print), the factor correlations (under oblimin or promax) for how distinct the factors are, and "Variance explained" for how much of the items' variation the factors carry. The preset retention rule is parallel analysis, applied to the eigenvalues before the factor tables are drawn.

Downloads and copying

The full set of seven outputs, subscription permitting, including the factor scores when requested.

When it refuses or warns

- Fewer than three items stops the run: "Select at least three items."
- A declared Taylor or replicate design is refused (the module supports none, sampling, frequency, and analytic weighting); the checks explain the supported set.
- A Heywood case (a communality at or above one) or a factor with no loading above threshold is flagged in the checks rather than smoothed over.

Module 14: Count, ordinal, and multinomial regression

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Four generalized linear models in one module, chosen by the outcome: Poisson regression for counts, negative binomial regression when counts are overdispersed (with a likelihood-ratio test of the two), ordinal logistic regression (the proportional-odds model, with a Brant test of the proportional-odds assumption) for ordered categories, and multinomial logistic regression for unordered categories against a baseline. An offset (exposure) variable serves the count models. Under a design, the fits carry design-based errors.

The data it needs

One "outcome", one to thirty "predictors", and an optional "offset" for the count models. The outcome's kind picks the family: non-negative counts for Poisson and negative binomial (three or more levels), an ordered categorical outcome for ordinal, an unordered one for multinomial. Worked example, a Poisson model with one binary predictor, entered as two rows, one per group: the first row has `events` 40, `exposure` 200, and `group` 1 (group A, a rate of 0.20 per unit of exposure); the second has `events` 25, `exposure` 200, and `group` 0 (group B, a rate of 0.125).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2. Choose "Module 14: Count, ordinal, and multinomial regression". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once an outcome and at least one predictor are in place, it lists one fitted model with coefficient intervals and the fit checks that match the model form.

3. Step 1, "Choose variables": the step description reads "The kind of outcome decides the model: a count, an ordered scale, or a category with several levels."

- "Outcome": click in the box and choose the outcome being modeled. It appears as a chip tagged with its type. The line under the box reads "The count, ordered scale, or multi-level category being modeled." - "Predictors": choose the variables the outcome is

modeled on. Pick at least one. Up to 20. - "Offset (exposure)" (optional): for count models, the exposure each count is a rate of, entered on the log scale. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the model": the step description reads "Choose the model that matches the outcome and how its standard errors are computed." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets agree on everything except the standard errors: the model-based form in "Essentials", the robust sandwich form in "Detailed". Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups:

- "Model": "Model", with "Poisson" (the preset), "Negative binomial", "Ordinal logistic", and "Multinomial logistic". The line under it reads "Poisson is the standard count model; the negative binomial adds room for wider spread. Ordinal fits an ordered scale, and multinomial fits a category with several levels." - "Precision": "Standard errors", with "Model-based" (the preset), "Robust (sandwich)", "Robust HC3", and "Cluster-robust". The line under it reads "Model-based errors assume the model has the variance right. Robust errors stay honest when it does not." Below it, "Confidence level": the level of the intervals printed for the coefficients: 90%, 95%, or 99%.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Baseline level: Not set · Random seed: 42". Inside:

- "Baseline level": for a multinomial model, the outcome level every other level is compared against. The alphabetically first level is the default. Type the level name or its code. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "purchases on 2 predictors · Essentials preset", or names the first thing still missing: "Choose an outcome and at least one predictor to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the model report, titled for the form you chose, for example "Poisson regression: purchases", with the coefficient table first.

Reading the output

The worked example prints a group coefficient of 0.47 (SE 0.25, p = 0.065), which exponentiates to an incidence rate ratio of 1.60: group A's event rate is 1.6 times group B's, per unit of exposure. Count-model coefficients are log rates, so exponentiate before speaking. For the ordinal model, coefficients are log cumulative odds, and the "Thresholds (cumulative logits)" table locates the category boundaries; read the Brant table before quoting a single odds ratio, since a rejected proportional-odds assumption means the effect

differs across boundaries. The "Observed vs model-expected count distribution" (count models) and "Observed vs average predicted category shares" (categorical models) tables are the fit check: systematic gaps there outrank any single coefficient.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- An outcome with two levels under the ordinal or multinomial family stops the run and points to Module 6: two levels are logistic regression's job.
- Negative values in a count outcome stop the run: counts must be non-negative.
- An offset declared for the ordinal or multinomial family stops the run: offsets belong to the count models.
- A baseline level that is not an observed level of the outcome stops the run, naming it.

Module 15: Mixed-effects models and repeated measures

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Three procedures for clustered and repeated data. A linear mixed model (LMM): fixed effects with random intercepts and an optional random slope for one or two grouping factors, crossed or nested, estimated by REML (preset) or ML, with Satterthwaite (preset) or Kenward-Roger degrees of freedom, profile intervals for the variance components, boundary-corrected likelihood-ratio tests for the random effects, predicted random effects per group (BLUPs), an intraclass correlation, and Nakagawa R squared.

A generalized linear mixed model (GLMM): binomial (logit) or Poisson (log) family with a single random intercept, fitted by adaptive Gauss-Hermite quadrature (25 points preset; the Laplace shortcut is not offered, because it did not meet the module's parity standard). And repeated-measures ANOVA: the classic within-subject analysis with Greenhouse-Geisser correction by default, its between-subject factor, cell means, and subject-level completeness enforced.

The data it needs

For the mixed models: one "outcome", fixed-effect "predictors", one or two "group" variables (the clustering factors), an optional continuous "slope" variable for a random slope, and a "trials" variable for binomial counts in the GLMM. For repeated-measures ANOVA: an "outcome", a "subject" identifier, one or two "within" factors, and an optional "between" factor. Interactions (mixed models only) are declared as pairs.

Worked example (the LMM, used in "Reading the output"): thirty pupils in six classrooms, five per classroom, with an end-of-term score and a treatment indicator (1 for the new method, 0 otherwise). `classroom` is 1 for the first five rows, 2 for the next five, and so on. `treat` by row: 0, 1, 0, 1, 0, 1, 0, 1, 0, 1, 0, 1, 1, 0, 1, 1, 0, 0, 1, 0, 0, 0, 1, 1, 0, 1, 1, 0, 0, 1. `score` by row: 60, 67, 63, 66, 61, 72, 66, 73, 69, 74, 64, 70, 71, 63, 72, 76, 70, 69, 75, 72, 65, 64, 72, 71, 67, 74, 75, 68, 67, 73.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".

2.

Choose "Module 15: Mixed-effects models and repeated measures". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: once an outcome is picked, it lists one fixed-effect table with tests and intervals, and the variance components for the random effects.

3. Step 1, "Choose variables": the step description reads "Pick the outcome, the fixed-effect predictors, and the grouping the random effects belong to."

- "Outcome": click in the box and choose the measure being modeled. It appears as a chip tagged with its type. - "Fixed-effect predictors": choose the predictors whose average effects the model estimates. Up to 20. - "Grouping variables": choose the clusters the cases belong to, such as respondents or sites: the random effects vary across these. Up to two. A grouping chip carries the group count instead of a type tag. - "Random-slope variable" (optional): the predictor whose effect is allowed to vary from group to group. - "Trials variable" (optional): for a two-level generalized model, the number of trials each case's outcome counts successes from. - "Subject variable" (optional): for the repeated-measures procedure, the case each set of repeated measures belongs to. - "Within-subject variables" (optional): for the repeated-measures procedure, the occasions or conditions of the repeated measures. - "Between-subject variable" (optional): for the repeated-measures procedure, a grouping that is constant within each subject. - "Interactions" (optional): build pairs of fixed-effect predictors with the "First variable" and "Second variable" boxes and "Add pair". Each pair appears as a removable chip, such as "visit_t:treat". The line under the control reads "Pairs of fixed-effect predictors whose combined effect the model also estimates. Add up to six pairs." - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the model": the step description reads "Choose the fitting procedure the outcome needs and how the groupings relate." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets agree on everything in this step. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings". The options sit in two groups:

- "Model": "Procedure", with "Linear mixed model (LMM)" (the preset), "Generalized linear mixed model (GLMM)", and "Repeated-measures ANOVA". The line under it reads "A linear mixed model for a continuous outcome, a generalized model for a two-level or count outcome, or the repeated-measures analysis of a balanced layout." Below it, "Outcome family", with "Binomial" (the preset) and "Poisson". The line under it reads "For the generalized model: a two-level outcome (logit link) or a count (log link)." - "Random effects": "Relation between two groupings", with "Crossed" (the preset) and "Nested". The line under it reads "When two grouping variables are picked: crossed when every level of one can occur with every level of the other, nested when the levels of one sit inside the other."

5.

Step 3, "Precision and tests": the step description reads "Choose the estimation criterion, how the fixed effects are tested, and the level of their intervals." The two presets differ here in one place: the denominator degrees of freedom, Satterthwaite in "Essentials" and Kenward-Roger in "Detailed". The options sit in two groups, with one switch below them:

- "Estimation": "Estimation criterion", with "REML" (the preset) and "ML". The line under it reads "Restricted maximum likelihood estimates the variance components without bias from the fixed effects. Full maximum likelihood suits comparing nested fixed effects." - "Tests": "Denominator degrees of freedom", with "Satterthwaite" (the preset) and "Kenward-Roger". The line under it reads "How the fixed-effect tests count their degrees of freedom: Satterthwaite or Kenward-Roger." Below it, "Confidence level": the level of the intervals printed for the fixed effects: 90%, 95%, or 99%. - "Profile intervals for variance components": a switch that starts on. The line under it reads "Print profile-likelihood intervals for the variance components of a linear mixed model."

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Quadrature points: 25 · Random seed: 42". Inside:

- "Quadrature points": how many adaptive quadrature points the generalized model integrates over, from 2 to 25. More points are slower and more exact. The value 1 is not offered: it did not meet the parity standard. Preset to 25. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "score within subj_id · Essentials preset", or names the first thing still missing: "Choose an outcome to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see the report titled for the procedure, for example "Linear mixed model", with the fixed-effect table first and the variance components after it.

Reading the output

On the worked example (score on treat, a random intercept per classroom, REML with Satterthwaite degrees of freedom): the random-effects table gives a classroom variance of 9.44 (SD 3.07) against a residual variance of 1.39 (SD 1.18), so classrooms differ materially before treatment is even considered; the boundary-corrected likelihood-ratio test for the classroom intercept is decisive (chi-square 37.32, p below 0.001).

The fixed-effects table gives an intercept of 65.96 (SE 1.29) and a treatment coefficient of 6.01 (SE 0.44, t = 13.67, p below 0.001, 95% interval 5.10 to 6.91): pupils on the new method score about six points higher. The ICC table reports an intraclass correlation of 0.41 (from the intercept-only companion fit), a marginal R squared of 0.46 (fixed effects alone) and a conditional R squared of 0.93 (fixed and random together); "Model fit" reports a log-likelihood of -55.92 and AIC 119.85 over 30 observations in 6 groups.

Read the random-effects table before the fixed effects: the grouping variance tells you how much of the outcome lives between clusters, and the intraclass correlation states it as a share. A grouping variance near zero means the clustering barely matters and a simpler model would tell the same story. The fixed-effects table reads like Module 5's, with the declared degrees-of-freedom method named. For repeated measures, the Greenhouse-Geisser corrected rows are the headline (sphericity is not assumed), and "Cell means" is the table to quote for the pattern over occasions. The likelihood-ratio tests for random effects are boundary-corrected, so their p-values already account for testing a variance at zero.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- No grouping factor for a mixed model, or no within factor or subject for repeated measures, stops the run naming the missing role.
- An interaction whose member is not a declared predictor stops the run: hierarchy is enforced, and the message names the component to add.
- A random slope with no grouping factor stops the run; more than two grouping or within factors are refused.
- A binomial GLMM outcome that is not coded 0/1 (and has no trials variable) stops the run, as does a Poisson GLMM outcome with negative or fractional values.
- Subjects without complete repeated measurements are excluded by the repeated-measures procedure, and the run reports how many: it analyzes complete subjects only, by design.
- Degenerate fits (a variance estimated at exactly zero) are reported in the checks rather than hidden.

Module 16: Survival analysis

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Time-to-event analysis: the Kaplan-Meier survival curve with a number-at-risk table and survival quantiles, log-rank or Peto-Peto comparison across a grouping variable (with Holm-adjusted pairwise tests), and Cox proportional-hazards models over covariates, with Efron tie handling (preset), optional robust (Lin-Wei) errors, an optional strata variable, and a proportional-hazards check from scaled Schoenfeld residuals. Under a declared design, the Cox model takes the design-based (survey Cox) construction and its tests become design-based Wald tests.

The data it needs

One "time" variable (time to event or to censoring; zero times must be genuine), one "event" indicator, and optionally a "group", up to twelve "covariates", and a "strata" variable. Worked example, eight cases with (time, event) reading (3, 1), (5, 0), (6, 1), (8, 1), (10, 0), (12, 1), (15, 0), (18, 1), where 1 marks the event and 0 marks censoring.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".
2. Choose "Module 16: Survival analysis". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once a time and an event indicator are in place, it lists one survival table and curve with intervals, and, when Cox covariates are picked, one Cox model table with hazard ratios.
3. Step 1, "Choose variables": the step description reads "Pick the time to the event and the column that says whether the event happened."

- "Time": click in the box and choose how long each case was followed, until the event or until follow-up ended. It appears as a chip tagged with its type. - "Event indicator": choose the column that says whether the event happened for each case; every other value counts as censored. - "Grouping variable" (optional): splits the cases into curves that are compared with the group test. - "Cox model covariates" (optional): when any are picked, a Cox proportional hazards model is fitted with them. - "Cox model strata" (optional): the Cox model gets its own baseline hazard within each stratum. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the analysis": the step description reads "Confirm the event coding, choose the interval method, and say whether a Cox model is fitted beside the curves." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets differ in three places: the group comparison test (log-rank in "Essentials", Peto-Peto in "Detailed"), the tied event times (Efron in "Essentials", Breslow in "Detailed"), and the robust standard errors (off in "Essentials", on in "Detailed"). Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in four groups, with three switches below them:

- "Event": "Event value": the value of the event indicator that means the event happened. Pick it from the levels found in the data, or leave the first level. Left unset, it shows "(none)". - "Intervals": "Interval type", with "Log-log" (the preset), "Log", and "Plain". The line under it reads "How the survival intervals are transformed: log-log, log, or plain." Below it, "Confidence level": the level of the survival intervals: 90%, 95%, or 99%. - "Group comparison": "Group comparison test", with "Log-rank" (the preset) and "Peto-Peto". The line under it reads "How the curves are compared: the log-rank test weights the whole follow-up evenly; the Peto-Peto form weights earlier differences more." - "Cox model": "Tied event times", with "Efron" (the preset) and "Breslow". The line under it reads "How the Cox model handles cases sharing an event time." - "Confirm the event value": a switch. The line under it reads "Required when the event coding is not plain 0 and 1: confirm the value above is the one that means the event happened." - "Confirm zero times are genuine": a switch. The line under it reads "Required when any follow-up time is zero: confirm those zeros are real times, not missing values." - "Robust standard errors": a switch. The line under it reads "Robust (Lin-Wei) standard errors for the Cox model coefficients."

6. There is no "Advanced settings" fold on this screen beyond the "Random seed", preset to 42, whose current value its summary line states.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "survival of surv_time by event · Essentials preset", or names the first thing still missing: "Choose a time and an event indicator to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Survival analysis" titled with your time variable, with the survival table first, then the curves and the group comparison.

Reading the output

The worked example's survival steps down at each event time: survival 0.875 at time 3, 0.729 at 6, 0.583 at 8, 0.389 at 12, and 0 at 18, so the median survival time (where survival crosses 0.5) is 12. Censorings shrink the risk set without stepping the survival, which the number-at-risk table makes visible. Cox coefficients exponentiate to hazard ratios: above one, the covariate raises the event rate at every time; below one, it lowers it. Read the proportional-hazards check before quoting a hazard ratio as a single number: a covariate whose Schoenfeld test objects has an effect that changes over time, and the honest report says so.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- A missing time or event variable stops the run, as does an empty complete-case set.
- An event indicator whose coding is not plain 0/1 blocks the run until the event value is set and confirmed: the message lists the values the indicator actually takes.
- Zero times block the run until confirmed as genuine.
- A group or stratum with no events, or a covariate with no variation among cases at risk, is flagged in the checks; the affected test or coefficient is not printed rather than printed wrong.

Module 17: Time series and forecasting

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

A one-variable-over-time analysis: the series is regularized onto an even time grid (multiple observations in a period combined by mean or sum; gaps filled by linear interpolation only when you approve), decomposed by STL into trend, seasonal, and remainder, tested for trend, and modeled by a candidate set of time-series models. The candidates are scored on a holdout (the last stretch of the series, kept back during fitting), the recommended model is named, and a forecast with intervals is produced for the horizon. An optional "group" variable runs the whole analysis separately per series.

The data it needs

One "time" variable and one "value" variable; the optional "group" splits into several series. Worked example: thirty-six monthly totals, January 2022 through December 2024, in two columns: `month`, the calendar month as a date (2022-01 for January 2022, 2022-02 for February, and so on through 2024-12), and `total`. A date column is offered as the time variable directly; a numeric time column (a period index, or fractional years such as 2022.00 for January 2022) works as well.

The totals, by year: 2022: 113, 117, 128, 137, 141, 148, 147, 146, 149, 146, 147, 148. 2023: 147, 157, 164, 169, 179, 183, 185, 183, 184, 183, 181, 184. 2024: 184, 192, 198, 208, 214, 221, 218, 221, 217, 219, 220, 218. The series climbs by about three a month with a yearly season (summers high, Januarys low). Your series needs at least two complete observations per series after regularization, and longer is materially better: seasonal reading needs several full seasons.

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin". This module reads its own time column, so a series file with one row per period is the natural input.
2. Choose "Module 17: Time series and forecasting". The setup is two numbered steps, with a "You will get" card at the side that updates as you choose: once a value is in place, it lists one decomposition of the series into trend, season, and remainder, and one forecast table and chart with intervals.

3. Step 1, "Choose variables": the step description reads "Pick the value measured over time and the time of each measurement."

- "Time": click in the box and choose when each observation was taken: a date column or a numeric time stamp. It appears as a chip tagged with its type. - "Value": choose the measure observed over time. - "Grouping variable" (optional): each level of the grouping gets its own series, decomposed and forecast separately. Its chip carries the group count instead of a type tag. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the series": the step description reads "Say how the observations are spaced and how many steps ahead to forecast." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom". The two presets differ in two places: "Detailed" switches "Fill gaps by interpolation" and "Robust decomposition" on. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings".

5. The basic options sit in two groups, with two switches below them:

- "Series": "Observations per season": the seasonal period of the series: 12 for monthly, 4 for quarterly, 52 for weekly. Leave empty to detect it from the time stamps. Below it, "Duplicate periods", with "Mean" (the preset) and "Sum". The line under it reads "When several observations fall in one period, combine them by their sum or their mean." - "Forecast": "Forecast horizon": how many periods ahead the forecast runs. Leave empty for two seasons, or ten periods for a series without seasons. - "Fill gaps by interpolation": a switch. The line under it reads "Missing periods are filled by linear interpolation, and every filled position is logged in the results." - "Robust decomposition": a switch. The line under it reads "A robust decomposition, less swayed by outlying observations."

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Holdout length: Not set · Random seed: 42". Inside:

- "Holdout length": the last periods are held back and forecast, so the forecast can be scored against what actually happened. Leave empty for the last fifth of the series or the last full season, whichever is longer. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "forecast of sales · Essentials preset", or names the first thing still missing: "Choose a value and a time to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Time series and forecasting" titled with your value, with the decomposition first, then the forecast table and chart.

Reading the output

On the worked example, the trend tests report a Mann-Kendall tau of 0.877 (p below 0.001) and Sen's slope of 3.000 per month (95% interval 2.73 to 3.22); the decomposition shows the trend rising from about 125 at the start of 2022 to about 224 at the end of 2024, with the seasonal component peaking in June (about +10.1) and troughing in January (about -14.4).

The "Candidate models" table flags the best ETS form and the best ARIMA as "Selected (AICc)" within their families (here ETS(A,A,N) at 156.26 and ARIMA(1,1,0)(0,0,0)[12] at 134.31). The module then compares those two on holdout MASE and builds the forecast from the better one: on the held-back last twelve months the ARIMA's MASE is 0.215 against the ETS form's 0.563 (its MAE is 7.74, against the drift benchmark's 7.79 and the seasonal naive's 35.92), so the ARIMA is the recommended model.

Had the better of the two failed to beat the seasonal-naive benchmark's own holdout MASE, the benchmark would have been issued as the recommended forecast instead, with a note saying so; the benchmarks are otherwise printed for comparison only. The "Forecast" table runs twenty-four months (the preset horizon, two seasons): the first, January 2025, is 219.28 with an 80 percent interval of 214.00 to 224.55 and a 95 percent interval of 211.21 to 227.35, and the points climb by about 2.94 a month.

Read in the order the output prints. The decomposition tells you what the series is made of (a rising trend with a yearly season reads exactly so). The forecast table gives the point forecast and intervals per period; intervals widen with distance, and a forecast whose interval spans your decision threshold is the module telling you the series cannot settle that question. When a group is declared, every table carries the group's suffix and one series is analyzed at a time.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Missing time or value variables stop the run; a series with fewer than two complete observations stops it as well.
- Gaps in the grid are reported, and with "Fill gaps by linear interpolation" off, a series too broken to regularize is refused rather than silently interpolated.
- Filled positions, when filling is on, are listed in the checks: a forecast leaning on interpolated stretches should be read with that in mind.

Module 18: Propensity score analysis

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

A treatment-effect estimate adjusted for measured confounding: a propensity model (logistic, on the declared covariates) estimates each case's chance of treatment; the effect is then estimated by inverse-probability weighting (the average treatment effect, ATE, or the effect on the treated, ATT) or by 1:1 nearest-neighbor matching on the logit of the score (ATT only). Balance is reported before and after adjustment as standardized mean differences with a love plot, intervals come from a bootstrap over the whole procedure (1,000 replications, seeded), and E-values summarize how strong an unmeasured confounder would have to be to explain the estimate away. Nothing runs until the causal assumptions are acknowledged, and the acknowledgment is printed with the results.

The data it needs

One binary "treatment", one "outcome" (continuous, binary as 0/1, or a non-negative count), and one to twenty-four "covariates" (the confounders). Worked example, eight cases with one binary covariate x, where treatment chances are 0.75 when x is 1 and 0.25 when x is 0: the cases (x, treated, outcome) read (1, 1, 10), (1, 1, 12), (1, 1, 11), (1, 0, 7), (0, 1, 6), (0, 0, 4), (0, 0, 5), (0, 0, 3).

The steps

1. Load data, and declare the design in "Survey design" if you have one. The "Survey design" card at the side of the setup repeats the current design in one line, with a "Set a design" or "Change" link back to that card. Until data is loaded, the action bar at the bottom reads "Load data to begin".
2. Choose "Module 18: Propensity score analysis". The setup is three numbered steps, with a "You will get" card at the side that updates as you choose: once an outcome and a treatment are in place, it lists one effect estimate with its interval, and the balance before and after adjustment, and propensity diagnostics.
3. Step 1, "Choose variables": the step description reads "Pick the outcome, the treatment, and the covariates the propensity model is fitted on."

- "Treatment": click in the box and choose the two-level variable saying who received the treatment. Its chip carries the group count instead of a type tag. - "Outcome": choose the measure the treatment effect is estimated on. It appears as a chip tagged with its type. - "Covariates": choose the background variables the propensity model adjusts for. Pick at least one. Up to 25. - A variable that cannot serve a role stays visible in the list, dimmed, with the reason stated under its name.

4. Step 2, "Set up the analysis": the step description reads "Choose the estimand and how the propensity scores are used, and confirm the assumptions the method rests on." Start from the "Preset" control, a segmented control with "Essentials", "Detailed", and "Custom": the two presets differ in the estimand (ATE in "Essentials", ATT in "Detailed") and in the Step 3 switches. Changing any basic option by hand moves the preset to "Custom". Presets never touch "Advanced settings". The options sit in three groups, with two switches below them:

- "Estimand": "Outcome type", with "Continuous", "Binary", and "Count", or left unset, shown as "(none)". The line under it reads "The kind of outcome. Leave unset to read it from the outcome variable itself." Below it, "Estimand", a segmented control with "ATE (average treatment effect)" (the preset) and "ATT (effect on the treated)". The line under it reads "The average effect over everyone, or the average effect over the treated." - "Method": "Method", with "Weighting (IPTW)" (the preset) and "Matching (1:1 nearest neighbour)". The line under it reads "Weighting reweights both groups to the target population. Matching pairs each treated case with the nearest untreated case." - "Treatment": "Treated value": the treatment level that counts as treated: its name as it appears in the data, or its code. The alphabetically last level is used when empty. - "Confirm the causal assumptions": a switch, and the run stays blocked until it is on. The line under it reads "The estimate is causal only if every confounder is among the covariates and every case could in principle receive either condition. Confirm to run." - "Proceed despite imbalance": a switch. The line under it reads "Shown only when weighting or matching leaves the groups imbalanced: confirm to print the effect anyway."

5. Step 3, "Refinements": the step description reads "Optional safeguards: trim extreme scores, cap the weights, or pair the estimate with an outcome model." The "Detailed" preset switches the last three of these on; "Essentials" leaves all four off. The four switches:

- "Stabilize the weights": the line under it reads "Stabilized weights for the average effect over everyone; they steady the weighting estimate." - "Trim extreme propensity scores": the line under it reads "Cases with propensity scores below 0.05 or above 0.95 are dropped before the effect is estimated, and the drop is logged." - "Cap the weights": the line under it reads "Weights are capped at their 1st and 99th percentiles, which steadies the estimate at some cost in bias. The capping is logged." - "Add an outcome model": the line under it reads "Fit an outcome model as well, so the estimate stays consistent if either the propensity model or the outcome model is right."

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values, for example "Bootstrap replications: 1000 · Random seed: 42". Inside:

- "Bootstrap replications": how many bootstrap samples draw the interval for the effect; preset to 1,000. - "Random seed": preset to 42. A changed value is marked "changed from default", with a "Reset to default" link beside it.

7. Press "Run analysis" in the action bar at the bottom of the panel. The bar recaps the setup as you work, for example "effect of treat on spend · Essentials preset", or names the first thing still missing: "Choose a treatment, an outcome, and at least one covariate to run", or, while the confirmation is off, "Confirm the causal assumptions to run". "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed. You should see "Propensity score analysis" titled with your treatment and outcome, with the effect estimate first, then the balance table.

Reading the output

The worked example's unadjusted difference is 5.0 (treated mean 9.75 against control 4.75), but treatment rides on x, and x also lifts the outcome. The weighted estimate is 3.0: weighting by the inverse chances rebuilds both arms to the same covariate mix, and three points is the effect that survives. Read the balance table before the estimate: standardized differences after adjustment should sit near zero; where they do not, the module asks for the imbalance acknowledgment and the estimate deserves a caveat in your report. The E-value translates residual doubt into a number: an unmeasured confounder would need associations at least that strong, with both treatment and outcome, to explain the estimate away.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Without the assumptions acknowledgment, nothing runs: the module prints the assumptions and waits. This is deliberate, and the acknowledgment is recorded in the run manifest.

-

A treatment with other than exactly two values stops the run and lists the values it found; a treatment coded other than 0/1 needs "Treated value" set to the name or code of the level that means treated.

- Matching with the estimand on ATE stops the run: matching estimates the effect on the treated.
- A binary outcome with values other than 0 and 1, or a count outcome with negative or fractional values, stops the run.
- No complete cases stops the run; extreme weights draw a check suggesting truncation or trimming, both of which are logged when used.

Module 19: Statistical Quality Control and Improvement

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Four procedures, chosen in one "Procedure" control. "Shewhart charts" plots control charts for a process in process order: X-bar and R, X-bar and S, and individuals and moving range (I-MR) charts for measurements, and p, np, c, u, and Laney p-prime and u-prime charts for counts, in Phase I (trial limits computed from the data) or Phase II (monitoring against frozen limits). "Run rules and chart performance" adds the Western Electric and Nelson run rules to the same charts, prints an average run length (ARL) disclosure for the rules in force, and offers estimation-effect-adjusted limits. "Measurement systems analysis" runs gauge studies on a measurement system: crossed and nested gauge R&R by the ANOVA method, with the Xbar-R average-and-range method as the legacy option. "Process capability" computes Cp, Cpk, Pp, and Ppk with intervals and one-sided lower bounds on the within and overall bases, by the normal model or by a Box-Cox, fitted-distribution, or nonparametric route, with the expected nonconforming rate in parts per million. Design of experiments, CUSUM and EWMA charts, multivariate charts, acceptance sampling, and the bias, linearity, stability, and attribute agreement studies within measurement systems analysis are not part of this module in this release.

The data it needs

One dataset in process order: the rows are the subgroups or the individual values in the order the process produced them. The roles are "Measurement" (the measured characteristic, or the event count for an attribute chart), "Subgroup" (the subgroup each row belongs to; leave it empty for individual values), "Subgroup size or units" (the number inspected for a p or np chart, or the inspection units for a u chart), and, for a gauge study, "Part" and "Operator". A gauge study needs balanced data: every operator measures every part the same number of times in a crossed study, or a distinct set of parts per operator in a nested study. Worked example for the charts and for capability: twelve fill weights in grams, in process order: 249.6, 249.9, 249.7, 250.2, 250.0, 250.4, 250.1, 250.6, 250.3, 250.8, 250.5, 251.0, against specification limits of 248.5 and 251.5.

The steps

1. Load data. The "Survey design" card at the side of this module's setup carries a fixed line instead of a design: "Survey designs do not apply: quality methods treat data as process or experimental data, not a survey sample." There is nothing to declare.
2.

Choose "Module 19: Statistical Quality Control and Improvement". The setup is five numbered steps, with a "You will get" card at the side listing what the chosen procedure returns: the chart with its limits and signals, the limits and estimates table, the signal log, the ARL disclosure, the gauge tables, or the capability indices, according to the procedure.

3. Step 1, "Data and order": the step description reads "The measurement, its subgroups, and the process order." Fill the roles the procedure needs, as above. The action bar at the bottom names what is still missing: "Load data to begin" with no data loaded, and, with data loaded, the measurement (for a gauge study, a study; for capability, the measurement, at least one specification limit, and the stability gate).

4. Step 2, "Chart": the step description reads "The chart type and whether limits are Phase I trial limits or frozen Phase II limits." Start from the "Preset" control, a segmented control with "Monitor a measurement", "Detect small shifts", "Gauge R&R (crossed ANOVA)", "Capability (normal model)", and "Custom": the presets set the procedure and its headline options together, and changing any basic option by hand moves the preset to "Custom". Other task combinations (defectives, defects, the nested gauge study, the non-normal capability route) are set by choosing the procedure and its options directly. The controls:

- "Procedure": "Shewhart charts" (the preset), "Run rules and chart performance", "Measurement systems analysis", or "Process capability". - "Chart type": "Auto" follows the registered default for the subgroup size (X-bar-R up to subgroup size ten, X-bar-S above ten, I-MR for individual values, p for counts with trials). Choose a type to override it: "X-bar and R", "X-bar and S", "Individuals and moving range", "p chart (proportion defective)", "np chart (number defective)", "c chart (defects)", "u chart (defects per unit)", "Laney p' chart", or "Laney u' chart". - "Phase": "Phase I (trial limits)" or "Phase II (frozen limits)". Phase I computes trial limits from all the subgroups you supply; the module never excludes a subgroup on its own. For Phase II, enter the "Frozen centre line" and "Frozen sigma" from your Phase I run, or the frozen limits themselves under "Advanced settings".

5. Step 3, "Rules and performance": the step description reads "Run rules, the ARL disclosure, and links to a gauge study."

- "Run rules": "Basic (3 sigma rule only)" (the preset), "Western Electric", or "Nelson". With "Nelson", the "Nelson tests" checkboxes choose among the eight tests, each named in plain words, from "One point beyond the 3 sigma limits" to "Eight points in a row with none in the inner zone". - "Shift for the ARL": the shift, in sigma units, at which the ARL is reported; preset to one. - "Adjusted limits": a switch offering estimation-effect-adjusted limits, which widen the Phase I limits so that, with the chosen probability, the conditional in-control ARL is at least the chosen fraction of nominal. Its settings (guarantee probability, ARL floor fraction, nominal ARL, adjustment estimator) sit under "Advanced settings". - "Linked gauge rating" and "Gauge study": the rating

and label of a gauge study run on the same characteristic; a poor rating raises a gauge caution on the chart.

6. Step 4, "Gauge study": the gauge design and method. The study runs on the "Measurement", "Part", and "Operator" variables from Step 1.

- "Gauge design": "Crossed" (every operator measures every part) or "Nested" (each operator measures a distinct set of parts, for destructive testing). Balanced data only; unbalanced data is refused with the reason. - "Gauge method": "ANOVA" (the preset), which carries the intervals described below, or "Xbar-R (average and range)", the labelled legacy method, which prints point estimates only. - "Tolerance": the specification width (upper limit minus lower limit). Enter it and the output carries percent of tolerance; leave it empty and percent of tolerance is absent, never silently zero.

7. Step 5, "Capability": specification limits, the stability gate and gauge link, the sigma basis, and the normal or non-normal route.

- "Lower specification limit" and "Upper specification limit": at least one is required. "Target" is optional; the target-based indices are refused when no target is entered. - The stability gate: switch on "Linked run in control" when a linked Phase I run on this characteristic was accepted as in control. Without it, switch on "Stability acknowledgment" to acknowledge that process stability is not demonstrated; with neither, the run refuses. The acknowledgment prints the sentence "Process stability not demonstrated; indices describe this sample only." on every capability output. - "Linked gauge %study variation": the linked gauge study's percent study variation; above thirty, every capability output carries the caution "Measurement variation is large; observed capability understates true capability." - "Within sigma estimator": pooled S by default with subgroups; R-bar, S-bar, or the moving-range estimators for individuals. The estimator and its degrees of freedom print with every index. - "Capability route": "Normal model" (the preset), "Box-Cox transformation", "Fitted distribution", or "Nonparametric percentiles". An Anderson-Darling check prints with the output and informs the choice; it never switches the route by itself. On the fitted route, "Fitted distribution" offers lognormal, Weibull, gamma, and Johnson, or fits all four and chooses by the Anderson-Darling statistic or by AIC, with the criterion printed and every family tabled. - "BCa bootstrap interval": a switch adding a bootstrap interval for one index, chosen at "BCa index" (Cpk, Ppk, or Cpm), with 2,000 replicates preset and the seed recorded in the manifest.

8. "Advanced settings" stays closed unless you need it. Its summary line states the current values. Inside: the frozen lower and upper limits, the known mean and known sigma standards, the moving-range span (preset to two), the individuals sigma estimator (mean moving range by default, median moving range as the robustness option), the S estimator, the adjusted-limits settings, the simulation runs behind a simulated ARL disclosure (preset to 100,000), the confidence level for the gauge and capability intervals (preset to 95 percent), and "Random seed", preset to 42.
9. Press "Run analysis" in the action bar. The bar recaps the setup with the measurement's name. "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed.

Reading the output

For the worked example on the I-MR chart, read the "Limits and estimates" table first: centre line 250.26, individuals limits 249.29 to 251.23 from a mean moving range of 0.36, and a moving-range upper limit of 1.19. Every one of the twelve points sits inside the limits, so the "Signal log" is empty under the basic rule set. The plotted-points table gives each value with its moving range and the limits in force, so any point can be checked by hand. The "Chart performance (ARL disclosure)" table states the in-control ARL of the rules in force (370.4 for the basic three sigma rule) and the ARL at the shift you chose; with a Nelson rule set the in-control figure is simulated at the run count under "Advanced settings", and the disclosure says so. Where the counts vary more than the binomial model allows, an overdispersion check fires and offers the Laney chart; the module never substitutes it silently.

A gauge run prints the ANOVA table, then the components table: repeatability, reproducibility, the part-to-part component, the combined gauge component, and total variation, each with its percent of study variation and, when a tolerance was entered, its percent of tolerance. The ANOVA method adds modified large-sample (MLS) intervals for the components that carry them; the output names the components whose intervals are not offered (the operator, part-by-operator interaction, and reproducibility components), so a component is never read as interval-backed when it is a point estimate only. The "Gauge study summary" adds the number of distinct categories, and prints the acceptance bands as what they are: conventions from the AIAG Measurement Systems Analysis manual (under ten percent study variation acceptable, ten to thirty percent marginal, over thirty percent poor; five or more distinct categories), displayed as conventions, not verdicts.

A capability run prints the indices on both bases. For the worked example against limits of 248.5 and 251.5: Cp 1.55 and Cpk 1.28 on the within basis, Pp 1.17 and Ppk 0.97 on the overall basis. Read the one-sided lower bound beside each index before the point value; the takeaway sentence does the same. The gap between the bases is the longer-term variation the within estimate does not see, and here it is the overall figure that falls below one. The output closes with the expected nonconforming rate in parts per million on each basis and the Anderson-Darling check.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- The np chart needs a constant subgroup size; varying sizes stop the run and point to the p chart. A subgroup of size one on the p chart is retained with full-width binomial limits, and on the X-bar-S chart it is plotted and excluded from the S-bar estimate; both behaviours print on the output.
- X-bar-R with a subgroup larger than ten draws a warning: X-bar-S is the default above ten.
- A gauge study with unbalanced data is refused with the reason, and the Xbar-R method is refused for the nested design, because that construction is crossed only.
- Capability without a linked in-control run and without the stability acknowledgment refuses; so do a missing specification limit, a lower limit not below the upper limit, the target-based indices without a target, and the nonparametric route with fewer than 300 values.
- Adjusted limits attach to the basic rule set only; combining them with the Western Electric or Nelson sets, and the two-sided dispersion adjustment, are refused with the reason.

Module 20: Epidemiologic Analysis

Access: free to run on screen during the promotion, through March 31, 2027. From April 1, 2027 this module needs the subscription to run. Downloads and "Copy Results" need the subscription at all times.

What it computes

Nine procedures, chosen in one "Procedure" control. "Measures of occurrence" computes a risk, an incidence rate, or a prevalence with its interval, from typed counts, from follow-up records, or under a declared survey design. "Standardization" adjusts rates and prevalences to a standard population, directly or indirectly (the standardized morbidity or incidence ratio, SMR or SIR). "Association and impact" computes the risk ratio, risk difference, odds ratio, number needed to treat, and the attributable fractions and risks for a two-by-two table, and the prevalence ratio, prevalence difference, and prevalence odds ratio under a declared survey design. "Stratified analysis and additive interaction" pools stratum tables by Mantel-Haenszel, tests homogeneity with the Tarone-corrected Breslow-Day test, and measures additive interaction (RERI, the attributable proportion, and the synergy index S). "Matched designs" analyzes matched pairs by McNemar and matched sets by conditional logistic regression. "Diagnostic accuracy" computes sensitivity, specificity, predictive values, likelihood ratios, the area under the ROC curve, and cutoff tables. "Agreement" computes kappa for two or several raters, the intraclass correlation in its six named forms, and Bland-Altman limits of agreement. "Outbreak investigation" draws the epidemic curve and the attack-rate table. "Life-table survival" builds the actuarial life table with case fatality. Every output opens with an orientation sentence stating what the rows and columns measure, and the interval method is named beside every adjusted or pooled figure.

The data it needs

This module is the one module besides the planner that opens with no dataset loaded, because counts are a primary way in. Three entry routes cover the release:

- **Typed counts.** The occurrence measures take their counts straight into the form: "Events", with "Persons at risk" for a risk or a prevalence, or "Person-time" for an incidence rate.
- **Dataset roles.** With a file loaded, the roles are "Outcome variable", "Exposure variable", "Age band", and "Subgroup", and, for person-time from records, "Entry date", "Exit date", "Event indicator", and "Birth date". Declaring a survey design in the "Survey design" card (Part 2) opens the design-based paths: prevalence under a design, design-based standardization by age band, and the survey path of association and impact.
-

Published counts. Where a procedure is built on a summary table of counts (a two-by-two table, one table per stratum, matched pairs, an accuracy table, a life table), enter counts as a small dataset with one row per cell. Summary-table entry as a form of its own, and a template download, are not part of the app in this release; a procedure whose table has not been collected stops and names the table it needs, as the last section of this chapter records.

Worked example for the occurrence measures: twenty-five events among 200 persons at risk, and the same twenty-five events over 1,850 person-years of follow-up.

The steps

1. Load data only if the run uses dataset roles or a declared design; the typed-count runs need none. Until a procedure's needs are met, the action bar reads "Choose a procedure and enter the counts the procedure needs."

2. Choose "Module 20: Epidemiologic Analysis". The setup is four numbered steps, with a "You will get" card at the side listing what the chosen procedure returns, from the chosen measures with their intervals to the orientation sentence, the impact measures with the attributable-fraction form stated, and the route to the survival module on the life-table output.

3. Step 1, "Procedure": the step description names all nine procedures. The variable pickers sit at the top of this step, as in every module; below them, the "Procedure" control carries the nine plain names listed above, and its line under the control summarizes what each one does, in the same order.

4. Step 2, "Data and design": the step description reads "Counts, person-time, or records; the study design for association and impact." Start from the "Preset" control at the top of this step, a segmented control with "Essentials", "Detailed", and "Custom": the presets set the headline methods named in Step 3 below, and changing any basic option by hand moves the preset to "Custom". The controls in this step follow the procedure:

- For occurrence: "Occurrence measure", a segmented control with "Risk", "Incidence rate", and "Prevalence"; then "Events", "Persons at risk", and "Person-time", as the measure requires. - For standardization: "Standardization", a segmented control with "Direct" and "Indirect". Under a declared design with an outcome and an age-band variable, the design-based path runs instead, and uses the built-in United States 2000 standard population when your age bands are its eleven bands. - For association and impact: "Study design", a segmented control with "Cohort", "Case-control", "Cross-sectional", "Person-time", "Survey", and "Diagnostic accuracy", chosen first because it decides which measures are offered. A case-control design offers the odds ratio only. "Survey" runs the design-based path on the outcome and exposure variables under the design declared in the "Survey design" card. - For diagnostic accuracy:

"Sampling form", a segmented control with "Cross-sectional sample" and "Two-gate sample", and "Prevalence for predictive values", the prevalence at which the Bayes predictive values are computed, as a number between zero and one.

5. Step 3, "Measures": the step description reads "The measures and interval methods."

- "Measures": the association and impact measures to compute, as checkboxes with plain names ("Risk ratio", "Risk difference", "Odds ratio", "Number needed to treat", "Attributable fraction in the exposed", "Attributable risk", "Population attributable fraction (Levin)", "Population attributable fraction (case-based)", "Population attributable risk"). The preset selection follows the study design. "Risk ratio is confounded" is a switch that flags the Levin population attributable fraction, which assumes no confounding. - Interval methods, each a select with its preset named first: "Proportion interval" (Wilson; Clopper-Pearson exact and Jeffreys are the alternatives), "Rate interval" (Garwood exact; Byar the approximation), "Adjusted rate interval" (Fay-Feuer gamma; Tiwari, Wald, and Dobson the options), "SMR interval" (Garwood exact; Byar), "Risk ratio interval" (Koopman score; Log-Wald the option), "Risk difference interval" (Newcombe hybrid score; Miettinen-Nurminen the option), "Odds ratio interval" (Automatic, which is exact conditional at small counts and Baptista-Pike otherwise; Woolf is the large-count option), "Rate ratio interval" (Exact conditional; Mid-p and Wald the options), and "Confidence level" (90, 95, or 99 percent; 95 is the preset). - Procedure controls in this step: "Interaction ratio scale" (risk ratios on the cohort scale, odds ratios on the case-control scale) and the "Bootstrap for S" switch for stratified analysis; "McNemar mid-p" for matched designs; "Paired comparison mid-p" for two-test accuracy comparisons; "Agreement proportion interval" and the "Bootstrap for kappa" switch for agreement; "Epidemic curve period" ("Onset hour", "Onset day", or "Onset week (Sunday start)") and "Attack-rate table order" for an outbreak.

6. "Advanced settings" stays closed unless you need it. Its summary line states the current values. Inside: the sparse-strata odds ratio switch and the adjusted risk ratio for stratified and impact work, the bootstrap replicate counts and seeds (2,000 replicates preset), the false-positive and false-negative costs and the cost-weighted prevalence behind the cost-weighted optimal cutoff, the design prevalence interval (Korn-Graubard by default, logit the option), the age band cuts and period cuts that split person-time by band (Lexis splitting), and "Random seed", preset to 42.

7. Step 4, "Run": the step description reads "Review and run." Press "Run analysis" in the action bar. "Reset", in the same bar, restores this module's defaults, and asks first when anything was changed.

Reading the output

Read the orientation sentence first; it fixes which level of which variable is the event and which group is the reference, and every figure after it follows that reading. For the worked example, the occurrence table gives a risk of 12.5 percent with a Wilson interval of 8.61 to

17.80 percent, and an incidence rate of 13.51 per 1,000 person-years with a Garwood exact interval of 8.75 to 19.95. Person-time from records prints the records, events, and person-years behind the rate, and the Lexis table splits events and person-time by age band or period when cuts are set. A prevalence under a declared design is design-based, and its note says what the undeclared path on the same data would have used instead.

Standardized figures carry the interval method's name beside every adjusted rate, and the indirect path prints the observed and expected counts behind the SMR or SIR. In stratified analysis, read the stratum-specific effects beside the crude and the Mantel-Haenszel pooled estimates, with the percent change and the ten percent heuristic labelled as a heuristic; the homogeneity tests print with their low power stated, and the additive-interaction block shows the four-category joint-exposure display behind RERI, the attributable proportion, and S. In matched pairs, the discordant counts drive everything: the McNemar exact and mid-p values and the matched odds ratio with its exact and mid-p intervals. Matched sets print conditional logistic coefficients as odds ratios, labelled a first-principles conditional-likelihood build.

For diagnostic accuracy, the sampling-form declaration and its gate outcome print with the results: under a two-gate sample the sample-prevalence predictive values are cleared, with the reason printed, and the predictive values at your stated prevalence are the default. A continuous test prints the area under the curve with its DeLong interval, the full cutoff table, and the Youden and cost-weighted optimal cutoffs as their full maximizing sets, with the cost direction stated in words. For agreement, kappa prints with its asymptotic interval; at low prevalence and small counts that interval can run narrow, and the methods text gives the measured coverage in that regime, with a warning on the output when the counts in hand fall in it. Percent agreement names its interval method, the six intraclass correlation forms are each named by design, type, and unit with their Shrout-Fleiss numbers, and Bland-Altman prints its difference convention (second minus first) with the bias and limits. An outbreak prints the epidemic curve as a count table with zero-filled periods over the observed range, and the attack-rate table carries both attack rates, the risk ratio with its Koopman interval, and Fisher's exact value for each candidate exposure, with the multiple-comparison caution printed with the table. The life table prints interval by interval (numbers at risk, deaths, withdrawals, the effective at-risk count, and the survival and hazard columns), survival read off the table with its Greenwood standard error, and case fatality with its interval method named. One route sentence stands on the life-table output, and this guide repeats it in full: "For Kaplan-Meier curves and Cox models, use the Survival Analysis module (Module 16)." The actuarial table is not a Kaplan-Meier estimate.

Downloads and copying

The full set of seven outputs, subscription permitting.

When it refuses or warns

- Events above the persons at risk, and person-time of zero, stop the occurrence measures before anything computes.
- Person-time from records needs both an entry date and an exit date; one without the other stops the run.
- A case-control design clears the risk measures and offers the odds ratio only, with the printed reason.
- The standardization design path with age bands that are not the built-in standard's bands stops and asks for standard population counts, one per age band.
- Procedures whose summary table has not been collected stop and name the table they need: a two-by-two summary table (or a person-time table, or the survey path) for association and impact away from the survey path, a stratified or joint-exposure table for stratified analysis, a paired table or matched sets for matched designs, an accuracy table, scores, or two-test inputs for diagnostic accuracy, ratings, an intraclass matrix, or Bland-Altman pairs for agreement, outbreak cases or candidate exposures for an outbreak, and a life table, case-fatality counts, or a survival proportion for life-table survival.
- Matched pairs with no discordant pairs refuse the matched odds ratio and say so. A kappa bootstrap whose failures exceed five percent of its replicates refuses the interval, with the failure count printed.

Saving and sharing your work

Analysis Studio is deliberately stateless between visits about your data: nothing is uploaded, and reloading the page starts with an empty "Data" card. What persists is your license key (in the browser, until you clear it from the "Forgot key" form) and your appearance choices. Everything else travels in files you save yourself:

1. Download the run's outputs while the session is open: the Word or Excel report for reading, the CSV for the tables, the run manifest for the audit trail, and the R script or SPSS syntax to reproduce the analysis elsewhere.
2. To continue later, keep your data file and your downloads together; re-running is load, declare the design, choose the module, set the same options and seed.
3. For weighting runs (Module 11 and its second procedure), the "New dataset" block under the results holds the merged data: download the new dataset during the session if you will need those weights outside the app, or press "Use this dataset now" to keep working with them in the app.

Sharing a result means sharing its downloads: the report files carry the tables, the takeaway sentences, the methods paragraph, and the checks, which is the complete defensible record. A number quoted without its methods paragraph and design declaration is a number without provenance; the manifest exists so you never have to reconstruct one.

Troubleshooting

What you see	What it means	What to do
"Load data first; the variable lists fill from your dataset."	The module panel opened before any data were loaded.	Load a file or press "Load the sample data", then reopen the module.
A variable does not appear in a picker	Its level does not fit the role (the hint under the picker lists the accepted levels).	Fix the variable's form in your file (labels for categories, numbers for measures), reload, and check the profile table's Level column.
A run stops under "Check the setup"	A required role is empty, a value is out of range, or a module's own requirement is unmet. The list names each problem.	Fix exactly what the list names, in order, and run again. Nothing is half-computed.
"Blocked:" with a button "Run unweighted, with the banner shown on the results"	The declared design is one this module cannot use (for example, analytic weights in a procedure with no design form).	Either change the design in "Survey design", or press the button to run unweighted; the result carries the banner "Unweighted analysis: design not applied", so it can never be quoted as design-based.
A download button is disabled with the licensed-set note	Downloads and copying need the subscription, including on free modules, including during the promotion.	Verify a key in the "Start Your Free Trial!" card, or work from the on-screen results.

What you see	What it means	What to do
"Run" is disabled on an advanced module with a license note	The promotion has ended (after March 31, 2027) and no key is verified.	Verify your key in the "Start Your Free Trial!" card; the form's entries are kept.
"Copy failed: the browser refused clipboard access."	The browser denied the clipboard write.	Allow clipboard permission for the page, or use a download instead.
The event or treated value warning appears	The module is about to model the last observed level as the event (logistic) or 1 as treated (propensity), and wants you to confirm that is what you mean.	Set "Event level (the outcome level modelled as the event)" or "Treated value" to the right value, or confirm, and rerun.
Results differ from a previous session	A different seed, a changed option, or a changed design declaration.	Compare the two run manifests; they record every input. Set the seed to 42 (or your recorded value) for simulation-based procedures.

Glossary

- **Analytic weight:** a weight saying each case's value is a mean built on that many observations. Contrast sampling weight.
- **ATE / ATT:** the average treatment effect over the whole population, or over the treated cases only (Module 18).
- **Average run length (ARL):** the average number of points a control chart plots before it signals; the in-control ARL is the false-alarm clock, and the ARL at a shift is the detection clock (Module 19).
- **Base weight:** the weight a case starts with before adjustment, usually the inverse of its selection probability (Module 11).
- **Calibration / raking:** adjusting weights so weighted totals match known population totals on chosen margins; raking fits several one-variable margins in turn, post-stratification fits one joint classification at once.
- **Checks:** the notes a module prints above its results when the data or the design changed how it computed, or something deserves a second look.
- **Complete cases:** the cases with no missing values among the variables an analysis uses.
- **Control chart:** a plot of a process measure in process order against a centre line and control limits computed from the process itself; points beyond the limits, and the runs the run rules name, signal that the process changed (Module 19).
- **Degrees of freedom (df):** the information budget behind a test or interval; designs and clustering reduce it, and the modules print which form they used.
- **Design-based:** computed under the declared survey design in "Survey design", with weights applied and standard errors from Taylor linearization or replicate weights.
- **Dummy coding:** representing a categorical predictor as indicator columns, one per level but the reference.
- **Effective sample size:** the size of a simple random sample that would carry the same precision as your weighted, clustered design.
- **Event:** the outcome level modeled as the one that happens in logistic regression and survival analysis.
- **Firth fit:** a penalized logistic fit that stays finite when a predictor separates the outcomes perfectly.

- **Gauge R&R:** a study of measurement variation: repeated measurements of parts by operators, split into repeatability (the same operator and part), reproducibility (between operators), and part-to-part variation (Module 19).
- **Hodges-Lehmann estimate:** a location estimate built from pairwise averages or differences, printed with the rank tests.
- **Kappa:** a chance-corrected agreement measure for two or more raters on categorical ratings: observed agreement beyond the agreement the marginal rates alone would produce, over the most that correction allows (Modules 19 and 20).
- **LMG / Shapley share:** a driver's share of explained variation averaged over every order the drivers could enter the model (Module 9).
- **Lonely PSU:** a stratum containing a single sampled cluster, needing a stated rule before a design variance can be computed.
- **Manifest:** the run manifest, the complete record of one run: module, inputs, options, seed, design, and checks.
- **Marginal effects:** a logistic effect restated on the probability scale, in percentage points, averaged over the cases.
- **Offset:** an exposure variable in a count model, so the model explains rates rather than raw counts.
- **Person-time:** the total time at risk summed over the people followed; the denominator of an incidence rate (Module 20).
- **Post-stratification:** calibration to one joint classification (Module 11).
- **Propensity score:** a case's modeled probability of receiving the treatment, given its covariates (Module 18); a pseudo-weight is built from a case's modeled probability of being in a nonprobability sample (Module 11, second procedure).
- **PSU:** primary sampling unit, the cluster sampled at the first stage of a survey design.
- **Pseudo-weight:** a weight for a nonprobability sample built from propensities against a reference probability sample.
- **Raking:** see Calibration.
- **Replicate weights:** extra weight columns, each a re-draw of the design, used to compute standard errors from the spread of replicate estimates.
- **Role:** a named slot in a module form (outcome, predictors, group) with its own accepted variable levels and counts.
- **Sampling weight:** a weight saying each case stands for that many population cases; the survey default.
-

Standard error (SE): the estimated standard deviation of a statistic across repeated samples; the width of every interval comes from it.

- **Standardization:** restating a rate or prevalence as it would stand in a standard population with a fixed age structure, directly from the study rates or indirectly from reference rates, so populations with different age mixes can be compared (Module 20).
- **Taylor linearization:** the standard design-based variance method, named for its first-order approximation.
- **Valid percent:** a percentage whose base excludes missing cases, beside percent, whose base includes them.