Bench sheet · EnGeL-SAMPL-POWER

Planning sample size: simulation-based power analysis for soil fauna and other ecological designs · version 1.1.0 (2026-11-20). Full protocol: EnGeL-SAMPL-POWER. Tick each step when done; write deviations in Notes.

Date	
Operator	
Plate ID / samples	
R and package versions	
Design, effect sizes, variances and number of simulations	

1. Set up

☐	No.	Step
☐	1	Load the packages and fix the random seed so the results can be reproduced:
☐	2	Build the design table:

2. Example 1: community composition (PERMANOVA)

☐	No.	Step
☐	3	Simulate a community matrix (plots × taxa) in which part of the taxa change in abundance with stand age. Each taxon has a baseline log-abundance; its response to age has a **fixed sign** (here 60% of taxa increase and 40% decrease) and a size scaled by `effect_size`; counts are drawn from a Poisson distribution with plot-level noise `noise_sd`:
!		**CRITICAL** Give each taxon's response a consistent sign. If the sign is drawn at random for each taxon, increases and decreases cancel out at the community level, PERMANOVA (which tests a shift of the community centroid) detects nothing at any effect size, and the power curve is wrongly flat at zero. A flat curve even at large effects is a sign of this error.
☐	4	Estimate power: simulate, compute Bray–Curtis dissimilarities, run PERMANOVA with the design factors and count significant results for the term of interest:
☐	5	Draw the power curve over a range of effect sizes and read off the MDE where the curve crosses the target power:

3. Example 2: a diversity index

☐	No.	Step
☐	6	For one value per plot (for example Shannon diversity or functional richness, standardised to mean 0 and SD 1 so that effects are in SD units), simulate the 2 × 2 design and test the age effect with a linear model:
!		**CRITICAL** Test the term that answers your question. With `y ~ land_use * age_class` and the default treatment contrasts, the coefficient `age_classmature` is the age effect **within the first land use only** (half the plots), not the overall age effect, and power is underestimated. For the overall (main) effect use the additive model above; to test whether the age effect differs between land uses, fit `y ~ land_use * age_class`, give `interaction_eff` a value and test the interaction term (`land_useLU2:age_classmature`).

4. Example 3: sampling methods, sub-samples and replicates

☐	No.	Step
☐	7	Pooling more sub-samples per plot makes the plot value less noisy. As a first approximation, the plot-level noise falls with the square root of the number of pooled sub-samples. Scale the noise of each method relative to a reference method:
!		**CRITICAL** The square-root rule assumes that sub-samples within a plot are as independent as plots are. Traps or cores close together often share microhabitat, so the real gain from pooling is smaller and the estimated power for methods with many sub-samples is optimistic. The number of plots, not the number of sub-samples, remains the sample size of the test.
☐	8	See how power grows with the number of replicate plots per cell, at a fixed effect size:

5. Example 4: repeated measures on the same units (tree-ring series, simr)

☐	No.	Step
☐	9	Build the unit-by-time data and a climate series:
☐	10	Create the model with chosen fixed effects, random-effect variances (a random intercept and a random climate slope per tree) and residual SD, then estimate power for the interaction:
☐	11	Repeat over a range of interaction sizes for the MDE, and over the number of trees per combination with `extend()` and `powerCurve()`:

Notes / deviations
