

Choosing species to study

To infer generalizations from samples, the rules are to replicate, to randomize and to circumscribe categories. Suppose we want to compare seedling density on upper versus lower slopes of a hill. The two categories, upper slopes versus lower slopes, need to be defined clearly. Then, within those categories, quadrats need to be placed at random (each possible location has an equal chance of being sampled) and replicated.

These inference rules forge the link between the sampling design of a study and the scope of generalizations that can be inferred from it. Boundaries on conclusions are set by the decision about what categories to compare. Legitimacy depends on the 'equal-chance-of-being-sampled' principle. The strength of conclusions can be assessed, because replication provides an estimate of variation within categories.

These inference rules are familiar. They were thoroughly assimilated, decades ago, into study designs for vegetation sampling and field experiments. But choosing species to study is not yet widely seen as a sampling problem of the same kind. In functional ecology and ecophysiology for example, respected papers have compared two or three species and interpreted them as representing categories, such as shade tolerant versus light demanding.

My own work has been far from perfect, so my purpose here is certainly not to denigrate past work. Rather, I hope to open a discussion about species-sampling practices for the future. Generalizations across categories of species could become much more reliable through improved species sampling.

The principle for comparing categories is clear. All species falling into each category are listed, and each species on a list should have an equal chance of being chosen as a replicate. Very few studies actually do this. Exclusion practices are one reason why. Some types of species that might commonly be excluded are:

- Species with too few individuals at a site
- Species that decline to germinate on cue
- Species with leaves inconveniently small for a gas exchange cuvette
- Rare and threatened species

- Species that are a long trek from a vehicle.

Exclusion practices are inevitable (although some do seem more virtuous than others). Regrettably, they are rarely discussed in publication. If exclusion practices were listed explicitly more often, literature reviewers could benefit greatly.

Phylogeny and choosing species

Species are often chosen on a phylogenetic basis. The three following phylogenetic designs illustrate some issues that arise:

Phylogenetically independent contrasts across two habitats

For each species chosen in one habitat, a related species is chosen in the other habitat, such that each pair is phylogenetically independent (PI) of other pairs within the study. The strength of this design is to produce a specified number of PI contrasts between the two habitats. A serious weakness is that it underestimates the overall difference between the species mixtures in the two habitats. This is because, in each habitat, a species is more likely to be included if it is closely related to species in the other habitat. Clades that only occur in one of the habitats are less likely to be included in the design.

Phylogenetically independent contrasts on a trait

These designs choose PI pairs of species having a minimum difference in some trait (e.g. threefold difference in seed mass). By pre-specifying a wide divergence, trait-contrast designs give good power in relation to the influence of the trait. A weakness of trait-contrast designs (and also of contrasts across habitat) is that only the distal PI contrasts can be used, not divergences deeper in the phylogenetic tree. Abstractly, this is because the species have not been chosen at random from the phylogeny. More concretely, imagine a clade in which most species have similar seed mass but one outlier species has a seed mass that is fivefold different from the others. This outlier will necessarily be chosen as one arm of the PI contrast within the clade. However, the average between this outlier and the other PI partner will

not give a fair estimate of the average seed mass of the clade.

Species within a genus

A very widely used design is to study a set of species within a genus (or other clade). The species might be compared across habitats or in relation to selected traits. The idea is that the basic similarity in morphology and in most traits within a genus will make for a better comparison. Working within a genus is seen as 'controlling' for unmeasured traits, at least partially. Actually though, the within-genus study is not a very good compromise. On the one hand, it is unreliable as a method of controlling for unmeasured traits, because all kinds of species trait can vary within a genus as well as between genera. On the other hand, by not replicating across several clades, the within-genus study fails to give any sense of whether patterns repeat themselves in multiple divergences.

In summary

To generalize effectively across species, we should pay more attention to how species are chosen for study. There is no single, simple design that works for all purposes, but some general recommendations can be made. First, make an explicit statement about the categories of species being compared, and the exclusion practices used. Second, be clear about the primary aim of the species-selection design. Is the aim to obtain maximum divergence on some trait that is a hypothetical causal variable? Or to maximize the total number of evolutionary divergences for a given number of species studied? Or to assess the consistency of the effect across different clades? Each of these would lead to a different design. Third, try to move at least a little in the direction of the ideal, which is to randomize from the full set of species within each category to be compared.

Mark Westoby

Dept of Biological Sciences, Macquarie University, Sydney, NSW 2109 Australia
e-mail: mwestoby@rna.bio.mq.edu.au

Published online: 14 October 2002