# A randomisation program to compare species-richness values

JEAN M. L. RICHARDSON and MIRIAM H. RICHARDS Department of Biological Sciences, Brock University, St. Catharines, Ontario, Canada

**Abstract.** 1. Comparisons of biodiversity estimates among sites or through time are hampered by a focus on using mean and variance estimates for diversity measures. These estimators depend on both sampling effort and on the abundances of organisms in communities, which makes comparison of communities possible only through the use of rarefaction curves that reduce all samples to the lowest sample size. However, comparing species richness among communities does not demand absolute estimates of species richness and statistical tests of similarity among communities are potentially more straightforward.

2. This paper presents a program that uses randomisation methods to robustly test for differences in species richness among samples. Simulated data are used to show that the analysis has acceptable type I error rates and sufficient power to detect violations of the null hypothesis. An analysis of published bee data collected in 4 years shows how both sample size and hierarchical structure in sample type are incorporated into the analysis.

3. The randomisation program is shown to be very robust to the presence of a dominant species, many rare species, and decreased sample size, giving quantitatively similar conclusions under all conditions. This method of testing for differences in biodiversity provides an important tool for researchers working on questions in community ecology and conservation biology.

**Key words.** Biodiversity, methodology, randomisation, species richness, statistical analysis.

#### Introduction

Ecologists are often interested in comparing biodiversity among two or more distinct communities separated in space or time, a statistical problem that has been mostly neglected (Colwell & Coddington, 1994; Mao & Lindsay, 2004). Despite the as-yet unresolved difficulties of estimating biodiversity (reviewed in Gotelli & Colwell, 2001; Magurran, 2004; McGill *et al.*, 2007), researchers comparing biodiversity among communities continue to focus on methods that require accurate and unbiased estimates of biodiversity in each community (Cao *et al.*, 2002; Buddle *et al.*, 2005; Cao *et al.*, 2007). Typically, comparisons among communities are accomplished by using rarefaction to generate estimates of the species richness of each community, estimates that have been standardised to the smallest available

Correspondence: Jean M. L. Richardson, Department of Biological Sciences, Brock University, 500 Glenridge Avenue, St. Catharines, Ontario, Canada L2S 3A1. E-mail: jrichardson@brocku.ca approach is necessary for biodiversity comparisons, because by collecting more samples (individuals), there is a greater chance of finding more different kinds of individuals; for example, in arctiid moths sampled from 21 sites that reflected a range of successional stages, the number of species observed per site was > 90% correlated with the number of individuals collected at the site (Hilt & Fiedler, 2005). The drawback of using rarefaction is that it limits the sample size to that of the smallest sample. Moreover, this sampling issue is not simply methodological: communities with greater abundance may also have greater diversity simply by chance. These difficulties of biased estimates are particularly prevalent in insect biodiversity studies because of the exceptionally large number of rare species typically encountered (Novotny & Basset, 2000; Hilt & Fiedler, 2005) and because the most thorough sampling of insect communities often fails to generate a species abundance curve that reveals its asymptote (Novotny & Basset, 2000; Longino et al., 2002; Buddle et al., 2005). Studies of insect biodiversity often require comparing species diversity in communities that

sample size (Magurran, 2004; McGill et al., 2007). This

Journal compilation © 2008 The Royal Entomological Society

differ markedly in abundance or which have been sampled unequally. Unfortunately, despite the extensive focus on finding biodiversity estimates that are unbiased and precise, all estimators currently available tend to be biased and have drawbacks that vary according to the data distribution (Hortal *et al.*, 2006; Cao *et al.*, 2007; McGill *et al.*, 2007).

While a good estimator of species richness may be beneficial, often our primary goal is to compare relative species richness of samples distributed in space or time, especially in a conservation context where the impact of anthropogenic factors may need to be assessed. Therefore, there is a need for a statistical method that improves our ability to compare relative species richnesses among samples; methods of estimating relative differences in species richness among assemblages may provide a more fruitful approach to solving this problem (Cao et al., 2007). This was the tack taken by Solow (1993) when he outlined a randomisation method to consider differences in biodiversity between two samples. He suggested an extension of this method to more than two samples that uses a measure of variance among the biodiversity of samples, a procedure subsequently carried out successfully by Bestelmeyer and Wiens (1996) and Wiens et al. (1996). Crist et al. (2003) described a randomisation method for estimating species diversity in the context of partitioning species diversity into  $\alpha$ ,  $\beta$ , and  $\gamma$  diversity. Randomisation was also suggested as an approach by Coleman (Coleman, 1981; Coleman et al., 1982) to test the null hypothesis of random placement of bird species onto islands, using species-area curves. Similarly, Koch (1987) considered sample size effects in species-richness measures by using the estimate of Fisher's  $\alpha$  from the data to calculate the probability of observing a certain number of species in a sample. Modern computers allow estimates of expected numbers of species without requiring that the data first be fit to a known distribution. Yet, despite the clear utility of randomisation methods for species diversity analysis, we found that no computer program exists to simply and intuitively allow a researcher to take advantage of these methods for testing hypotheses regarding differences in species richness among communities. Such a method would test the null hypothesis that communities were drawn from the same regional species pool, that is, that observed values of species richness are not improbable under the null hypothesis that  $\gamma$ -diversity does not differ. The alternative hypothesis is that communities differ in species richness and that the observed differences among communities are unlikely under the assumption that both communities were drawing from the same regional species pool. This paper presents a simple program that allows users to run readily randomisation tests for comparisons of species richness among samples.

### Materials and methods

The program code was written and compiled in Microsoft Visual C++ 5.0 (©1997, Microsoft Corporation, Redmond, CA, USA). Instructions for use of the program, the program source code, and the compiled executable program are freely available from the first author's website (http://www.brocku.ca/researchers/jean\_richardson/spp\_richness). The program runs in a DOS-

window and has no special requirements; any personal computer with a DOS-based operating system can run the program.

As an example of the analysis method, consider a study designed to test differences in species richness among three different forest types, with five replicate samples collected for each type. The appropriate null expectation for no difference in species richness among samples or groups is that all communities were drawn from the same underlying population (i.e. that the three forest types have communities drawn at random from the same regional species pool, implying that  $\gamma$ -diversity is equal for all samples) and that individuals occur at random and with equal probability in every sample. To test for differences among groups, we need to estimate how likely it is that the observed data would occur if individuals from the common species pool are distributed at random among forest types (Manly, 1997). This requires an estimate of the distribution of species-richness values among sites under the assumption of random placement. To generate this distribution under the null model (all communities were derived from the same regional species pool), we pool all individuals of every species into one large group (our best estimate of the common species pool) and then randomly re-assign individuals to samples, with the constraint that total abundance in each sample and for each species remains fixed at observed values. This differs from Coleman's curve (Coleman, 1981) because we do not need to assume that all sites have the same amount of resource. In this way, we can test whether any observed differences in species richness among sites can be explained solely by differences in the abundance of organisms at different sites, or whether additional factors, such as habitat quality, behavioural differences in habitat choice, etc., might also be at work.

Once all individuals have been randomly re-assigned to a sample, species richness is calculated for each sample and also for each group. This generates a set of 'pseudo-values' of species richness under the null hypothesis. The whole procedure is repeated 10 000 times to generate a frequency distribution of expected species richness per sample and per group (each group composed of > 1 sample), based on the null hypothesis of no difference in regional species pools being true. In the standard manner, the observed values of species richness can then be compared to those expected under the null hypothesis, and if the observed species richness is less than the lowest 2.5% or greater than the highest 2.5% of the randomisation estimates (for any of the samples in the analysis), we conclude that the underlying species pools likely differ (Manly, 1997). Note that because all samples are tested simultaneously, and these samples are not independent, we reject the null hypothesis if any one of the samples has a species-richness value that falls in the upper or lower tail of the probability density curve for whatever  $\alpha$  value we choose.

The technique was assessed using both simulated and real data. Simulated data were generated based on a log-normal curve; while the log-normal (and log-series) curve inadequately describes the fit of species-abundance data (Williamson & Gaston, 2005), we use the log-normal curve because of its historical dominance in the literature (e.g. Baltanas, 1992; Cao *et al.*, 2002). The number of species present, *y*, in each of five octaves to either side of the modal octave is given by:

© 2008 The Authors

**Table 1.** Results of replicate runs of simulated data sets using the randomisation program to test for species-richness differences between sites that do not differ (simulation 1), that differ in abundance but not diversity (simulations 2, 3), and that differ in diversity (simulations 4–6). Proportion of times null is rejected is based on 5000 independent and randomly drawn samples (from a regional pool of 200 species and 258 840 individuals) for each simulation. Sites were presumed to be sampled exhaustively and site differences in species richness for each of the 5000 samples were tested using the randomisation program provided here, with 5000 iterations per test. Note that rejection rates for each site are given for information purposes, but that the two tests in each simulation are not independent. Therefore, rejection of the null hypothesis for either site leads to rejection of the overall null hypothesis that the two sites share the same regional species pool.

		Site N*	Relative richness	$H_0^{*}$	Proportion of tests that reject $H_0^{\dagger}$
Simulation 1	Site 1	500	100%	True	0.0501
	Site 2	500	100%		0.0487
Simulation 2	Site 1	500	100%	True	0.0692
	Site 2	200	100%		0.0628
Simulation 3	Site 1	500	100%	True	0.0612
	Site 2	300	100%		0.0684
Simulation 4	Site 1	500	100%	False	0.3476
	Site 2	500	50%		1.000
Simulation 5	Site 1	500	100%	False	0.1016
	Site 2	500	80%		0.9410
Simulation 6	Site 1	500	100%	False	0.1296
	Site 2	200	50%		1.000

N =population size;  $N_0 =$  null hypothesis.

†Using a 2-tailed  $\alpha = 0.05$ .

$$y = y_0 e^{-a^2 R^2} \tag{1}$$

with  $y_0$ , the modal number of species = 50, a = 0.2, and R = the number of octaves away from the mode (Preston, 1948, 1962; Cao et al., 2002). The number of individuals per species within each octave was determined as 2<sup>R</sup>y<sub>0</sub>; this number was rounded to an integer, used as the median number of individuals per species in each octave, and then the range of species' abundances for each octave were determined as non-overlapping values given these median values. The rarest species had an abundance range of one to three individuals. The abundance for each of 200 species (distributed across R as dictated by equation 1) was drawn at random from within the range of the appropriate octave. This produced a simulated 'universe' of 200 species and 258 870 individuals. Simulated sites were then created by drawing individuals at random from this simulated universe. To test the behaviour of the statistical model, two types of simulations were carried out: (i) all species could persist in both sites and sites had n = 500, n = 300, or n = 200 individuals; and (ii) only a subset (50% or 80%, on average) of species could persist in one of the two sites and sites had n = 500, n = 300, or n = 200 individuals. Species site use was determined by randomly drawing a number between 0 and 1; species with a number above the proportion (0.5 or 0.8, according to site) could be present in only one site. Thus, in the first scenario, any differences that arise in measured species richness reflect only differences in abundance and our analysis should fail to reject the null hypothesis. In the second scenario, sites differ in species richness and our analysis should reject the null hypothesis of no difference in species richness among sites. For six different combinations of site size and species richness, 5000 replicate random sites were drawn and the number of times the program found a species richness significantly different from expected based on the null hypothesis (i.e. the number of times the

observed species richness fell within the 2.5% region at either end of the probability density curve) was determined. For these simulations, the number of iterations run by the program to generate the probability density curve for each random pair of sites drawn was reduced to 5000 (from 10 000) to make computing time feasible. These simulated data provide a check of the type I and type II error rates generated by the program.

Data published by Grixti and Packer (2006) are used to demonstrate use of the program on a real data set. The data consist of bee species abundances for a field site in southern Ontario that was sampled in 1968 and 1969 ('early') and then again in 2002 and 2003 ('late'). The abundance of each species in each year is given in Table 1 of Grixti and Packer (2006); a total of 20 221 individual bees, representing 165 species, were collected and identified in the four years. The sample abundance curve of their data differed significantly from both log-normal and log-series distributions (log-normal:  $X^2 = 31.02$ , P = 0.003; log-series:  $X^2 = 1270$ , P < 0.001; Fig. 1). Grixti and Packer (2006) were interested in comparing the diversity estimates of bees in early versus late periods, and did this by comparing rarefaction curves (estimated through re-sampling; Gotelli & Entsminger, 2004) for each period. A drawback of this method is that the curves can only be compared at points with equal sample sizes, so the information present in larger samples is lost. This difficulty does not exist with the method presented here, and the randomisation program was used to make the same comparison of species-richness values in the early versus late period. The input to the program includes the summed abundances of each species over all four years. These abundances are used to calculate expected values of species richness in each year, and in each period, under the null hypothesis that the variance in species richness between years is due only to random sampling variation (including variation in sampling effort).

© 2008 The Authors

Journal compilation © 2008 The Royal Entomological Society, Insect Conservation and Diversity, 1, 135–141



**Fig. 1.** Species-abundance curve for all four years of bee data published in Grixti & Packer (2006), along with expected number of species using a log-normal or log-series curve fit. Expected values were generated using PRESTONDISTR and FISHERDIST in the package VEGAN (Oksanen *et al.*, 2007) designed for R (R Development Core Team, 2007). Observed values differed significantly from both the log-normal and the log-series fit.

### Results

Simulated data indicate that the program behaves as expected when the null hypothesis is true (Table 1). Recalling that the null hypothesis is that all sites included in the analysis were drawn from the same regional species pool, we reject the null hypothesis if any one of the samples in our study falls outside the middle 95% of the distribution of predicted sample species-richness values (Fig. 2). The type I error rate when sites were drawn from the same regional species pool and were equal in size was the expected value of 0.05, indicating that the statistical test is behaving correctly (Table 1). A difference in abundance between sites (even though individuals were drawn from the same species-richness pool, i.e. the null hypothesis was true), led to a slightly inflated type I error rate of between 0.06 and 0.07 (Table 1).

Power of the analysis when sites differed in diversity was high (Table 1). When 50% of species could use one site, the null hypothesis was rejected 100% of the time, and when 20% of species were limited to site 1 only, the null hypothesis was correctly rejected 94% of the time. This result was robust to a decreased number of individuals present in the less diverse site (simulation 6) as might be expected to occur in nature (Table 1).

The randomisation program (with 10 000 iterations) was next used to analyse the field data collected by Grixti and Packer (2006). Analysis of the early versus late sampling periods clearly rejects the null hypothesis of no difference in species richness. Plotting observed versus expected species richness for each period allows us to assess further why the null hypothesis is rejected. In particular, it reveals that in the late sampling period, the observed species richness was similar to that expected, whereas in the early period it was much lower than expected under the null hypothesis (Fig. 3).



**Fig. 2.** An example of the frequency distribution of 5000 expected species-richness values generated by the program for a sample of 200 individuals. This is a representative example of the hypothesis-testing procedure, using simulated data in which one sample comes from a population with only 70% species presence. The filled bars in the tails of the histogram show the bottom and top 2.5% of the frequency distribution. The arrow points to the observed species richness for this sample, which clearly falls below the critical value for P < 0.05.

The analysis also included expected values for each year (two early and two late; Grixti & Packer, 2006 present no analyses of separate years). The data reveal a significant difference among years, the years 1968, 1969 and 2002 clearly having significantly lower species richness than expected based on the null hypothesis of no difference among the four years (P < 0.0001 for each year; even the most conservative Bonferroni adjustment recognises these individual tests as statistically significant). Conversely, species richness in 2003 was very similar to that expected (Fig. 4A). The difference in observed species richness between early and late periods was interpreted by Grixti and Packer (2006) as being due to an increase in species richness in the late period. Our analysis suggests instead that species richness was higher in 2003 than in any other year. Looking at the raw data indicates that 36 of the 165 species were observed only in 2003.

To consider the possibility that the above analysis was biased by large numbers of rare species, the analysis was repeated, removing all species whose combined abundance in the 4 years was < 10 individuals. This reduced the total number of species to 100 and the total number of individuals to 20 007, a 39% decrease in species but only a 1% decrease in total abundance. The general pattern of observed and expected values was similar in this analysis, although 2003 observed species richness was now significantly lower than expected based on random (Fig. 4B). This is not surprising as 31 of the 36 species found only in 2003 were rare.

Further consideration of the data revealed that one species, *Lasioglossum* (*Dialictus*) *imitatum* (Smith), accounted for 28.5% of individuals in the 4 years combined and 58% of those collected in the early period. Furthermore, this species was 7.2 and 4.8 times more abundant than the second most abundant species in 1968 and 1969, respectively, whereas the most abundant

© 2008 The Authors



**Fig. 3.** Expected and observed species-richness values for early and late periods of Grixti and Packer's bee data (2006). Expected values were calculated based on 10 000 iterations of the randomisation program that randomly assorted species back into the four sample years (constraining sample abundance to observed values) and then calculated species richness for the 2 years in the early period combined and the 2 years in the late period combined (i.e. a species present in both 1968 and 1969 would only be counted once).

species [Andrena (Taeniandrena) wilkella (Kirby) and Ceratina (Zadontomerus) calcarata Robertson, respectively] in 2002 and 2003 were only 1.6 and 1.1 times as abundant, respectively as the second most abundant species in those years. Therefore, a third analysis was carried out in which *L. imitatum* was removed from the data set, in case it was having undue influence on the randomisation results. The input data for this analysis had 14 452 individuals and 164 species. The expected species richness for 1968 and 1969 was decreased when *L. imitatum* was excluded because the abundance for those years were reduced by 59% and 56%, respectively, but the overall pattern of results among years remained similar (Fig. 4A).

Variation between runs in the model was very low. In six replicate runs of the program using the full data set, the coefficient of variation in mean species richness for each of the four sample years ranged from 0.0185% to 0.0249%.

## Discussion

The randomisation program presented in this paper provides a robust analysis of differences in species richness compared to a null hypothesis of no difference among samples. Analysis of simple data collected from populations which were generated from known distributions shows that the randomisation test performs as expected both when the null hypothesis is true and when it is false. While the statistic had slightly inflated type I error rates with differences in abundance, they were still reasonably close to the  $\alpha = 0.05$  criterion used and were consistent regardless of abundance differences; thus, a simple adjustment of the rejection criterion can be used where it is important that the type I error rate not exceed 0.05.



**Fig. 4.** Expected and observed species-richness values for each year of Grixti and Packer's (2006) data separately (n = 6290, 3494, 4293, 6144 for 1968, 1969, 2002, 2003, respectively). (A) Expected and observed values for analyses with all individuals included, as well as expected values when the most abundant species, *L. imitatum*, was removed from the analysis, reducing the total number of individuals by 28.5%. Note that *L. imitatum* was present in all years. (B) Expected and observed species-richness values for analysis with species whose total abundance in the 4 years was < 10 individuals removed, reducing the total number of species from 165 to 100 (39%) but reducing the total number of individuals only from 20 221 to 200 007 (1%).

An important distinction between this and previous methods for analysis of species richness is that this analysis does not attempt to estimate the true species-richness value. Rather, the goal was simply to test the null hypothesis that any difference we observe in species richness is attributable entirely to sample size differences. If any one of the predicted distributions within the analysis does not include the observed value for that sample, then the conclusion can be reached that abundance differences among samples are not sufficient to account for the observed differences in species richness and that there are other ecological processes contributing to the observed biodiversity patterns.

# © 2008 The Authors

Journal compilation © 2008 The Royal Entomological Society, Insect Conservation and Diversity, 1, 135-141

Simulations 2 and 3 (Table 1) support the appropriateness of this statistical test by showing that type I error rates are not tied to site abundance.

The incorporation of a second level of grouping further allows analysis of hierarchically organised data sets, such as that provided by Grixti and Packer (2006), in which both year and period could be analysed simultaneously, quickly revealing that differences between periods actually reflected a strong effect of year. Grixti and Packer (2006) used bootstrapping to estimate confidence intervals for species richness in each period and for the difference in species richness between periods. While this is a common method for estimating confidence intervals around species-richness estimates, it is known to be less accurate than other methods of species-richness estimation, particularly the Chao2 and ICE estimators (O'Hara, 2005) and to consistently underestimate species richness (Chiarucci et al., 2003; Hortal et al., 2006). Grixti and Packer's (2006) analysis, which kept data for each period isolated, led the authors to conclude that species richness had increased in the recent period. By incorporating both year and sampling period into the analysis, not only can a hypothesis test regarding species richness differences among samples be generated, but an assessment of where those differences lie can also be carried out. Thus, this new analysis allowed us to infer that recent species-richness values are as expected in the late period of the Grixti and Packer (2006) data, whereas species richness in the early period was reduced, and that this reduction might have been a statistical artefact resulting from the dominance of a single species, L. imitatum, in the early period (Lennon et al., 2004; Magurran, 2004).

The randomisation method of this program was extremely robust to the distribution of input data. Analysis that included a dominant species (accounting for nearly 30% of the abundance) reached the same conclusion as analysis that removed this species. The model was also robust to the presence of rare species; analysis with all rare species (abundance < 10 individuals) removed reached the same conclusions as the analysis that included these data.

Not surprisingly, predicted species-richness values based on random placement of individuals are typically greater than those observed. This is not a necessary artefact of the data analysis; data in which many species are present in low abundance in one site only, lead to observed species-richness values that are greater than expected by chance when individuals are randomly reassigned (data not shown). In nature, however, individuals are not distributed completely at random with respect to species; rather, individuals of the same species are more likely to be near one another, as demanded by sexual reproduction and the dispersal distances and patterns of offspring. Nonetheless, relative differences in species richness will be unaffected if the assumption can be made that these ecological effects on distribution are similar for all species included in the study.

An advantage of our program is that it allows the user to test hypotheses regarding differences in species richness even for samples with unequal sampling effort. For example, ecologists often wish to ascertain whether areas of lower density are also less species rich. But species-richness estimates at low-density sites will necessarily be based on small sample sizes. Methods using rarefaction curves to compare species richness between sites are limited by the number of individuals collected in low-density sites, whereas our randomisation method can accommodate differences in sample sizes and sampling effort without loss of power and with only a minimal increase in type I error rates. We contend that the randomisation program presented here can provide a robust, flexible, statistically powerful, and easy-to-execute analysis of species-richness differences among communities.

#### Acknowledgements

Thanks to Amy Rutgers-Kelly for troubleshooting the program as it was built. This research was funded by the Department of Biological Sciences, Brock University, the National Science and Engineering Council of Canada (DG# 261587-03 to JMLR; DG# 222883-03 to MHR), the Canadian Foundation for Innovation and the Ontario Innovation Trust (Project #9369 to JMLR).

#### References

- Baltanas, A. (1992) On the use of some methods for the estimation of species richness. *Oikos*, 65, 484–492.
- Bestelmeyer, B.T. & Wiens, J.A. (1996) The effects of land use on the structure of ground-foraging ant communities in the Argentine Chaco. *Ecological Applications*, 6, 1225–1240.
- Buddle, C.M., Beguin, J., Bolduc, E., Mercado, A., Sackett, T.E., Selby, R.D., Varady-Szabo, H. & Zeran, R., M. (2005) The importance and use of taxon sampling curves for comparative biodiversity research with forest arthropod assemblages. *Canadian Entomologist*, **137**, 120–127.
- Cao, Y., Hawkins, C.P., Larsen, D.P. & Van Sickle, J. (2007) Effects of sample standardization on mean species detectabilities and estimates of relative differences in species richness among assemblages. *American Naturalist*, **170**, 381–395.
- Cao, Y., Williams, D.D. & Larsen, D.P. (2002) Comparison of ecological communities: the problem of sample representativeness. *Ecological Monographs*, **72**, 41–56.
- Chiarucci, A., Enright, N.J., Perry, G.L.W., Miller, B.P. & Lamont, B.B. (2003) Performance of nonparametric species richness estimators in a high diversity plant community. *Diversity and Distributions*, 9, 283–295.
- Coleman, B.D. (1981) On random placement and species-area relations. *Mathematical Biosciences*, 54, 191–215.
- Coleman, B.D., Mares, M.A., Willig, M.R. & Hsieh, Y.-H. (1982) Randomness, area, and species richness. *Ecology*, 63, 1121–1133.
- Colwell, R.K. & Coddington, J.A. (1994) Estimating terrestrial biodiversity through extrapolation. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 345, 101–118.
- Crist, T.O., Veech, J.A., Gering, J.C. & Summerville, K.S. (2003) Partitioning species diversity across landscapes and regions: a hierarchical analysis of alpha, beta, and gamma diversity. *American Naturalist*, 162, 734–743.
- Gotelli, N.J. & Colwell, R.K. (2001) Quantifying biodiversity: procedures and pitfalls in the measurement and comparison of species richness. *Ecology Letters*, 4, 379–391.
- Gotelli, N.J. & Entsminger, G. (2004) ECOSIM: Null Models Software for Ecology. Version 7. Acquired Intelligence Inc. & Kesey-Bear. Jericho, Vermont. Available from URL: http://garyentsminger.com/ecosim/.
- Grixti, J.C. & Packer, L. (2006) Changes in the bee fauna (Hymenoptera:

© 2008 The Authors

Apoidea) of an old field site in southern Ontario, revisited after 34 years. *Canadian Entomologist*, **138**, 147–164.

- Hilt, N. & Fiedler, K. (2005) Diversity and composition of Arctiidae moth ensembles along a successional gradient in the Ecuadorian Andes. *Diversity and Distributions*, **11**, 387–398.
- Hortal, J., Borges, P.A. & Gaspar, C. (2006) Evaluating the performance of species richness estimators: sensitivity to sample grain size. *Journal of Animal Ecology*, **75**, 274–287.
- Koch, C.F. (1987) Prediction of sample size effects on the measured temporal and geographic distribution patterns of species. *Paleobiology*, 13, 100–107.
- Lennon, J.J., Koleff, P., Greenwood, J.J.D. & Gaston, K.J. (2004) Contribution of rarity and commonness to patterns of species richness. *Ecology Letters*, 7, 81–87.
- Longino, J.T., Coddington, J. & Colwell, R.K. (2002) The ant fauna of a tropical rain forest: estimating species richness three different ways. *Ecology* **83**, 689–702.
- Magurran, A.E. (2004) *Measuring Biological Diversity*. Blackwell Publishing, Malden, Massachusetts.
- Manly, B.F.J. (1997) Randomization, Bootstrap and Monte Carlo Methods in Biology, 2nd edition. Chapman & Hall, Boca Raton, Florida.
- Mao, C.X. & Lindsay, B.G. (2004) Estimating the number of classes in multiple populations: a geometric analysis. *The Canadian Journal of Statistics*, **32**, 303–314.
- McGill, B.J., Etienne, R.S., Gray, J.S., Alonso, D., Anderson, M.J., Benecha, H.K., Dornelas, M., Enquist, B.J., Green, J.L., He, F., Hurlbert,

A.H., Magurran, A.E., Marquet, P.A., Maurer, B.A., Ostling, A., Soykan, C.U., Ugland, K.I. & White, E.P. (2007) Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. *Ecology Letters*, **10**, 995–1015.

- Novotny, V. & Basset, Y. (2000) Rare species in communities of tropical insect herbivores: pondering the mystery of singletons. *Oikos*, 89, 564–572.
- O'Hara, R.B. (2005) Species richness estimators: how many species can dance on the head of a pin? *Journal of Animal Ecology*, **74**, 375–386.
- Oksanen, J., Kindt, R., Legendre, P., O'Hara, B. & Stevens, M.H.H. (2007) VEGAN: Community Ecology Package. R package. version 1.8-8. Available from URL: http://cran.r-project.org/, http://r-forge.rproject.org/projects/vegan/.
- Preston, F.W. (1948) The commonness, and rarity, of species. *Ecology*, **29**, 254–283.
- Preston, F.W. (1962) The canonical distribution of commonness and rarity: part 1. *Ecology*, **43**, 185–215.
- R Development Core Team (2007) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-9000551-07-0, Available from URL: http:// www.R-project.org.
- Solow, A.R. (1993) A simple test for change in community structure. *Journal of Animal Ecology*, **62**, 191–193.

Accepted 24 April 2008