

Chapter 8. Fossil J-curves

Although it might seem an odd thing to do, a biologist conducting a biosurvey could not only count the individuals per species collected, he or she could also count the number of species per genus, erecting a histogram for such counts. In this case, the first column would indicate the number of genera that had just one species in the area under survey. The second column would indicate the number of genera with two species, and so on. Indeed, for the sake of examining the patterns thus produced, it would not be necessary for the biologist to confine the count to the area being surveyed. The same exercise could be carried out over a much wider area with the help of a field guide or synopsis. Not stopping at species per genus, the biologist could also count the number of genera per family. Amazingly, the shapes seen in the earlier sampling efforts involving individuals per species would reemerge.

This chapter provides mathematical support for the proposition that taxonomic abundances are essentially inherited, so to speak, from species abundances in past epochs. It then demonstrates the presence of the logistic-J distribution in these taxonomic distributions by means of the Kolmogoroff-Smirnov test applied to some 55 sets of taxonomic data. Although this can only be described as a preliminary study, it strongly supports the idea that taxonomic distributions are, almost literally, fossil J-curves.

8.1 Background of the problem

Up to this point the main focus of this monograph has been present-day communities, sets of species, with each species being a set of individuals. But a consideration of communities in the distant past leads us to a) a taxonomic structure that goes beyond these two categories and b) a concept of “community” that expands to continental and even global scales.

If we take only the traditional taxonomic levels (omitting suborders or superfamilies, for example), we find a hierarchy of sets, each set being subdivided by its predecessor.

individual < species < genus < family < order < class < phylum < kingdom

We have already seen that when the species in a community are sorted into abundance categories a J-shaped curve emerges. Indeed, the phenomenon extends not only to consecutive hierarchies, but any pair at all: If we choose two different taxonomic levels from the hierarchy and examine the distribution of taxa at the higher level into abundances of taxa at the lower level, the same shape of histogram emerges. For example, if we distribute all the known families of North American herpetofauna (Conant & Collins 1981) into those having just 1 genus, then those with 2, then 3, and so on, we obtain Figure 8.1.

The shape of the distribution in Figure 8.1 is very similar to the species-abundance distributions we have been examining up to now. This is not the first time the phenomenon has been

observed. Our early hero, C. B. Williams (1964), the British biologist, summarized his research at Rothampstead over the 1940s and 1950s as follows:

“One general principle, or mathematical pattern, has several times been discussed. In almost every classification that has been proposed the number of genera with a single species is greater than the number with two, the number with two greater than with three, and so on. If one plots such a classification in the form of a frequency distribution we get a hollow curve, not unlike a hyperbola, which immediately recalls the similar pattern already discussed for the relative abundance of species.”




Figure 8.1. Numbers of families containing a given number of genera
(One family with 35 genera occurs off-page.)

Over the last century, others have noticed the “hollow” shape of taxonomic distributions, notably Willis and Yule (1922), Udney (1924), Burlando (1990), and Chu and Adami (1999). In most of the cited research a power law is proposed, compared with a few datasets and justified in terms of branching processes in (abstract) taxonomic trees. The simplest power law is from the first authors cited. One may apply it to species and genera, as follows.

$$G(k) \propto k^{-\gamma}, \tag{i}$$

Here, α represents a proportionality relation, k is a taxonomic abundance in species, $G(k)$ is the number of genera having k species and γ is a constant lying in the interval [1.4, 1.6]. We note in passing that a power law with exponent near unity is very close to a logistic-J distribution, the underlying hyperbola having the formula $k^{-\gamma}$, where $\gamma = 1$. The exponent 1 is just outside the interval specified in equation (i). Perhaps the difference is due to the absence of a logistic limit in the power law formula. .

Chu & Adami (1999) propose a somewhat more general scheme called the Galton-Watson process that produces a sequence of generations (whether of species, genera or other taxa) governed by a simple probability distribution p that for each taxon and each possible number k of “offspring” (a lower taxon) gives the probability $p(k)$ of that many offspring at the next “generation”, presumably a rather long stretch of time. The authors compare the results of the process under slightly different values of the average number of offspring, from a pure power law to a negative exponential distribution. The results are compared with one dataset drawn from the paleontological record using a Kolmogoroff-Smirnov test and the fits are about as good as one might expect from a distribution that resembles the logistic-J. In any event, the authors are under the same impression as others in the field have been, that a single set of data, whether from the present or the past, makes a reasonable test for a hypothesis.

Returning for a moment to the community of limited extent, the biologist would notice that some species that were known to exist in the surrounding area had yet to be observed in the community itself -- and perhaps a genus or two, as well. Once it was understood that the phenomenon was rooted deeply in the past, it would also be understood that the distribution was best not limited to the area under study, but would include the general area, the region and even extend to continental or global scales. Over such extents of time and space, genera themselves have undergone a stochastic vibration of sorts, losing species and gaining them by extinction and speciation, respectively. Given the fact that early genera may become later families, one expects to see the phenomenon extend right up the taxonomic tree to the point of emergence of DNA-based life itself. We do.

8.2 Ancient communities and trees of descent

The setting for an exploration of the influence of past logistic-J communities on present ones is a *tree of descent* in which every organism that ever lived (including ones in the present) are each represented by one or more nodes. Directed edges, one entering each node except the root (r) node, trace descent from parent to offspring. Each node may have any finite number of directed edges leaving it, from zero (death) to one (continued existence or death accompanying a single offspring) or greater than one (multiple progeny). Such a structure, apart from its interpretation, is called a directed tree in the graph-theoretic sense (Gross and Yellen 2004).

Successive nodes in the tree are joined by a downward directed arc that represents the time of one reproductive cycle for the organism represented by the node. Since reproductive cycles may vary

from minutes to years, depending on the organism so represented, some arcs may be much longer than others. In any case, the tree is divided into contemporaneous levels and, for sexually reproducing species, males are omitted. These structural features would be sufficient to bring our tree of descent into a reasonable alignment with the actual tree without burdening it with more detail than is necessary.

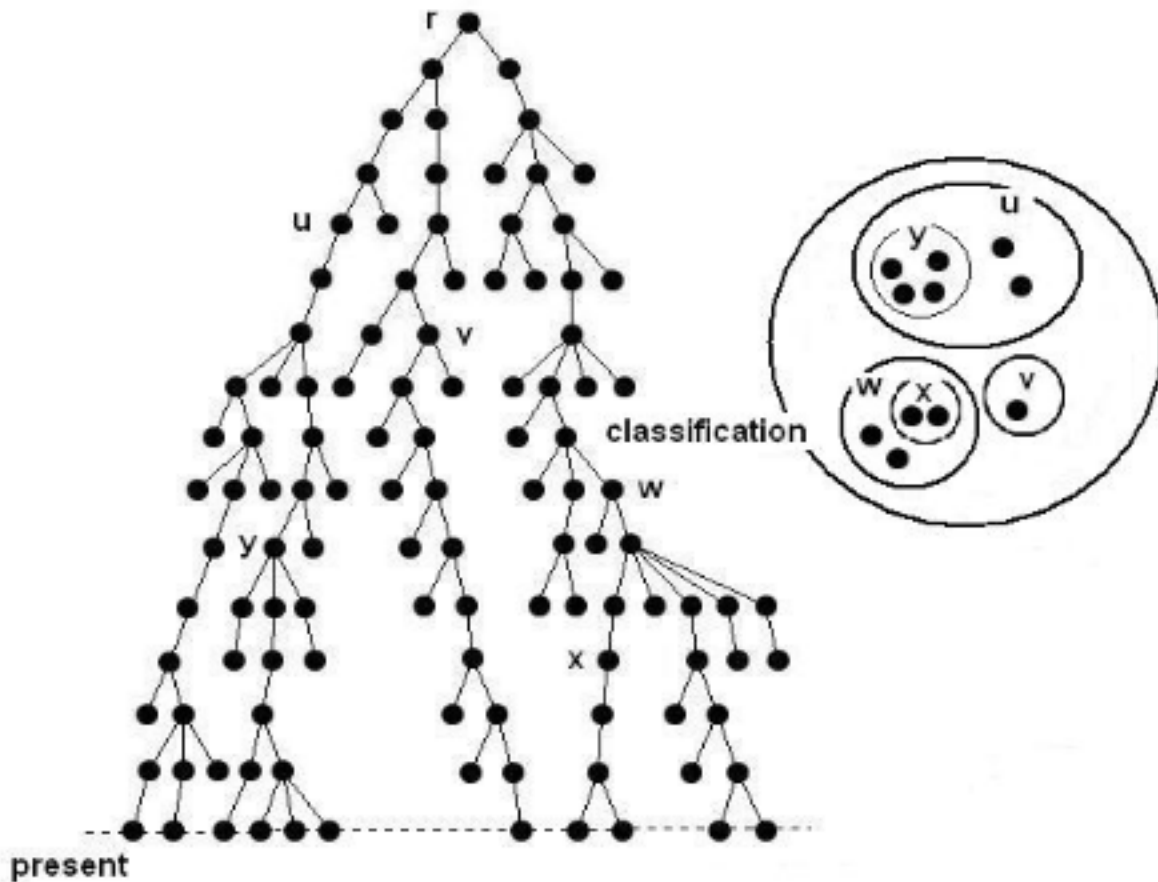


Figure 8.2 a tree of descent and its associated classification

Figure 8.2 shows a small, simplified example of a tree of descent in order to illustrate the notation employed in the argument to follow. To be the size of a “real” tree of descent, it would of course have not dozens of nodes but billions of trillions. In this humble version, however, we can observe that each node (organism) v in the tree has a pendant branch that includes all descendants of that organism. This branch at v will be a tree in its own right, called T_v . If u is a node of T_v , then T_u is a branch of T_v , as well as of T . If neither node is a descendant of the other, the branches T_u and T_v will be entirely disjoint, as shown in the figure. We refer to the set of all nodes at the bottom (present time) of T_u as U and the set of all such nodes at the bottom of T_v as V . In what follows, the terms “branch” and “subtree” will have exactly the same meaning, with the second term emphasizing that a branch is actually a tree in its own right. Figure 8.2 also

displays the set structure associated with the tree. The sets U and V appear there, as well as the descendants of another branch, W.

To introduce taxa, I will assume that the nodes u, v, w, x, and y are sites of speciation, whether relatively sudden, through gradual phyletic change, or even genome acquisition (Williamson 2009). In each case of gradual change, we select a node for which the process may be said to be complete. The corresponding sets have capital letter labels. They clearly have a nested structure, with inner circles representing species, while outer circles represent genera. We note that a genus may itself continue as a species, even though a taxonomist at the bottom of the tree would classify the two isolated individuals within the set U as members of a second species.

8.2.1 Branch weights in trees of descent

Let the node u now represent the common ancestor of all the species belonging to a particular genus. The subtree T_u descends to the present day with each species represented by a current population, as in Figure 8.2. Given (supernatural) knowledge of each site of speciation within this branch, we would find a set of nodes scattered more or less randomly throughout the branch, some later in time and some earlier. The argument to follow makes the reasonable assumption that since heavier branches have more nodes, they will harbour proportionally more speciation events, other things being equal.

Starting with an arbitrary node u of the tree of descent at a given time t, let T_u be the subtree at u and define the function $W(t)$ as the sum,

$$W(t) = \sum N(k), \tag{ii}$$

where $N(k)$ denotes the total number of descendants of u at time k, the summation being taken from $k = t$ to $t + b$. Here the parameter b is arbitrary but fixed over the duration of the argument. As the time variable increases, the portion of T_u spanned by the b units of time amounts to a large, moving block of nodes and the weight W to a rolling sum. The only important feature of this large portion of T_u is its potential size.

When the time t increases by unity, the first term in the sum (ii) is dropped and a new term added at the end. The new value,

$$W(t+1) = W(t) + N(t+b) - N(t),$$

yields a net increase of

$$N(t+b) - N(t).$$

at time t+b.

As shown in Section 6.2.1, the population behaviour of species in a weakly stochastic community may be characterized by a binomial distribution (with zero mean) that governs the probability of an increase (or decrease) by k units. Owing to the symmetry of the binomial distribution, it may be concluded that the probability of a community increasing by k individuals will equal the probability of a decrease by the same amount. Although the resulting expressions may be complicated, it will be seen that the change in the quantity $W(t)$ induced by the addition of $N(t+b)$ is governed by the same expression that determines the change in $W(t)$ induced by the subtraction of $N(t)$. The rolling sum W therefore follows the logistic-J distribution.

The forgoing argument implies a picture of how the rolling sum $W(t)$ behaves over time. It could be described as a stochastic orbit, but with the possibility of large changes, as well as small ones. Most often this quantity will change only a little as time progresses. Occasionally, however it will undergo greater changes and, rarely, rather large ones. Although the argument might well be made into a theorem, I am content to leave it as a conjecture. (See Section 9.

Conjecture: The weights $W(t)$ follow the logistic-J distribution.

Given the truth of this conjecture, the way would then be open to the following view of evolution in a logistic-J world: Imagine the tree of descent as it may have appeared at time $t = 100$ million years ago, say. Let the baseline b take the value of 100 thousand years and examine the weights $W(t)$ of all rolling subtree-blocks. An examination of a particular rolling subtree over the stretch of time from t to $t + 100,000$ yrs might reveal a few speciation events. The larger the subtree block, other things being equal, the more such events it would be expected to produce. A similar assumption underlies the analysis of Chu and Adami (1999), where the longer a taxon has been in existence, the more likely it is to spawn sub-taxa.

From the perspective of the taxonomy of current life forms on the planet Earth, one might look back to a particular speciation event u (if one only could) and identify it as the birth of a new species. Tracing further back in time from u , there is only one way to go until one encounters another, older speciation event v , one that today's taxonomists would now identify as the birth of a particular genus. And so it goes, up the tree and up the taxonomic hierarchy as one went, to the birth of a family, then an order, and so on.

8.3 Presence of the logistic-J distribution in taxonomic data

The statistical tests applied to the taxonomic data used here not only evaluated the logistic-J distribution as a descriptor of taxonomic abundances, but were used to evaluate scale-restricted and less formal sources of taxonomic data for bias in the resulting statistical scores. Table 8.2 lists the source of information for each taxonomic group used in the study, as well as the extent of geographic and taxonomic coverage. In the table, I have indicated taxa by one-letter codes (e.g., "s" = species, "g" = genus, etc.) This makes it possible to calculate the number of tests available

| Taxonomic group & range | # | Geographic coverage | Reference |
|------------------------------------|----------|----------------------------|-----------------------------|
| Plantae [f to c] | 3 | global | ITIS 2002b |
| Gymnospermatophyta [s to f] | 3 | North America* | Kartesz 1994 |
| Pteridophyta [s to f] | 3 | North America* | Kartesz 1994 |
| Animalia [o to p] | 3 | global | ITIS 2002c |
| Mammalia [s to o] | 6 | North America | Hall 1981 |
| Reptilia & Amphibia [s to o] | 6 | North America | Conant & Collins 1991 |
| Aves [s to o] | 6 | global | Sibley & Moore, 1990 |
| Pisces [g to o] | 3 | global | Nelson 1984 |
| Testudines [s to f] | 3 | global | Iverson 1992 |
| Insecta [f to o] | 1 | North America | Borror & White, 1970 |
| Arachnida [s to f] | 3 | E. North America | Caston 1972 |
| Collembola [s to f] | 3 | global | Bellinger et al., 1996-2003 |
| Pogonophora [g to f] | 1 | global | Ivanov 1063 |
| Ciliophora [g to o] | 3 | global | Lynn 2007 |
| Foraminifera 1988 [g to f] | 1 | global | Loeblich & Tappan, 1988 |
| Fungi [f to c] | 3 | global | ITIS 2002a |
| Bacteria [s to p] | 3 | global | Holt et al. 1994 |
| Life [c to p] | 1 | global | Wilson et al. 1973 |
| total datasets: 55 | | * including Greenland | |

Table 8.2. Sources of taxonomic information

from the data. For example, [s to g] means “species to genus” and only one set of data, species within genera, is available. On the other hand, [s to f] or “species to families” makes three datasets available: species within genera, species within families, and genera within families. I have placed the number of datasets so derived immediately after this notation.

For some groups it was not possible to obtain the coverage desired. For example in Class Pisces (Nelson 1994) the ostensible coverage was [s to o]. But for many genera, this global synopsis

gave no precise figure for the number of species therein, providing only an estimate, such as “about 60.”

Another example, the Fungi (ITIS 2002a), illustrates the use of online data in this inquiry. The Integrated Taxonomic Information System (ITIS) was still under construction at the time of this particular research. It yielded only [f to c] coverage, owing to having many genera being absent from its lists. At higher levels, with no apparent omissions, the data were assumed to be complete. The remaining ITIS-based tests involved plants and animals. In the latter case, all data pertaining to Phylum Ciliophora were omitted, as these do not properly belong to Kingdom Animalia. More recently, an authoritative taxonomic table relating to Phylum Ciliophora was found (Lynn 2007).

The Bacterial data, drawn mainly from *Bergey's Manual* (Holt et al. 1994), listed species within genera and the latter within “Groups.” I interpreted this word as synonymous with phyla or divisions, as in standard treatments on bacteria, the names being largely the same. (See Margulis and Schwartz, 1982, for a summary.) At the same time the Manual gives no taxonomic abundance of species for 21 genera, stating that the determination of the species within these genera involves tests that are too advanced for the Manual.

The last taxonomic group, namely “life” itself, involved an attempt to compile all the classes within the approximately 92 phyla of life on the Earth. In addition to the source listed in Table 8.2, I used the general reference (Margulis and Schwartz, 1982) and several general web-based resources such as the Tree of Life Web Project (Tol 2001) and the UK Systematics Forum (UKSF 2004). Phyla not covered in the foregoing sources included Sipunculan worms (Cutler 1994) and several protistan phyla (Margulis, Corliss et al. 1990).

8.3.1 The test method

In the metastudy of species abundance data, I used the chi square test, but for this investigation I switched to the Kolmogoroff-Smirnov test, partly because of its alleged sensitivity, but also to demonstrate the confirmatory approach with a completely different goodness-of-fit test. When a multiple-data test of the sort described here produces pass-fail ratios that are as close to the ones predicted by the Kolmogoroff-smirnov (See Table C2 in Appendix C), there is little room for any other distribution to play a role. To put it another way: any other distribution that succeeds to a greater extent with this taxonomic abundance data must nevertheless be very similar to the logistic-J distribution.

Data from all the sources in Table 8.2 were put into histogram form for each taxonomic abundance pair. In the case of species within families, for example, I counted the number of species in each family for the group in question, creating a standard histogram (as in Figure 8.1). The data so organized were subjected to the Kolmogoroff-smirnov test for goodness-of-fit. As explained in Section 4.4, the test statistic depends not on maximum differences.

$$D = \max \{|F(k) - G(k)|; k = 1, 2, \dots n\}$$

In the context of this chapter, $F(k)$ will be the number of higher taxa that contain k instances of a lower taxon, as in the number of families, for example, that contain k species. The function $G(k)$ will be the corresponding value of the logistic-J function. How close is the predicted number of higher taxa to the actual number? Is the maximum difference great enough to cause D to exceed the critical value of the test? In the context of a comparison between a taxonomic histogram and the logistic-J distribution, the critical value amounts to a kind of envelope around the logistic-J distribution. If the histogram values all fall inside this envelope, the comparison is given a pass.

In the K-S test, there is a critical value for each size of sample and for each level of significance. Respective scores achieved by each taxonomic abundance histogram were compared with the critical value for 95%, 90%, and 80% levels of significance. A histogram that passed at the 95% level of significance might well fail at 90% or 80% since in those cases, the test is more lenient in regard to rejection. If one wants only a small probability of being wrong in rejecting a particular histogram the critical value is higher; if the difference D exceeds even this value, the histogram is even more likely not to fit. It follows then, that one would expect in many cases to find acceptance (“pass”) at one level of significance, while finding rejection (“fail”) at a lower one.

The respective scores for each taxonomic abundance histogram were compared with critical values in order to decide whether the data “passed” or “failed” the test. If the histograms collectively follow the underlying Kolmogoroff-Smirnov distribution, one would expect the number passing at the 95% level to be 95% of the 55 histograms thus tested. Similarly, 90% should pass at the 90% significance level and 80% at the 80% level.

Scores were also converted into equivalent scores at a nominal, but fixed, sample size of 25. Each score was normalized by applying a correction factor of $\sqrt{N/5}$, where N is the sample size of the dataset. Having normalized scores made it possible to compare scores evaluated at different sample sizes.

Three questions arise naturally in this context:

1. Does the common origin of taxonomic histograms within a source group influence their independence?
2. Does it make any difference whether one examines continental-scale data, as opposed to global-scale data?
3. Does it make any difference that field manuals, rather than formal synopses were sometimes used?

To answer these questions, the normalized test scores were subjected to simple tests that compared average scores and variances between the data in question and the overall results. The next section addresses these questions first

8.3.2 Results of the study

Over all 55 taxonomic datasets the (normalized) average K-S score was 0.166 and the average (normalized) variance was 0.0064. The question of statistical independence within source groups was addressed by examining whether the scores within groups tended to be clustered, the most meaningful measure in this context. The contributions to overall variance from within groups ranged from small values to numbers in excess of the overall variance of 0.0064. One may check this important observation by examining the results of the study, as listed in Appendix C2. Variances within the data presented in the three largest tables are displayed in Table 8.3. In the first two cases the variance handily exceeds the overall variance and in the third table it come out slightly less.

| table | variance |
|--------|----------|
| first | 0.014 |
| second | 0.023 |
| fifth | 0.006 |

Table 8.3. Variances within the three largest K-S score tables

In short, the numbers give no support to the proposition that the common origin of data within a group reduced its overall variation. The assumption of statistical independence of the datasets therefore seems justified.

The question of whether continental-scale data differed from global-scale data in influencing the outcome of the main test was addressed by comparing the continental-scale score with the overall average. Table 8.4 lists the 18 scores derived from continental-scale data. The average score of 0.169 barely exceeded the overall score of 0.166 for the general form of the distribution. In neither case could it be said that using continental-scale data compromised the results by producing lower scores. No further testing of this issue was deemed necessary.

Results for field guides versus synopses were not so clear cut. Although for some of the groups covered, such as birds, field guides would be taxonomically complete at a continental scale, others, such as Insecta (Borer and White 1970), could not not be, owing to the vast number of insect species in comparison with the number of pages in a typical book (or website). As it happened, field guides came in with an average score of 0.117, somewhat below the overall average score of 0.166. Although it could be said that using field guides favoured the theory

under test, the small number (11) of datasets derived from field guides may have resulted in coincidentally low numbers. A brief test of this possibility involved the score for a field guide not used in the study (Whittaker, 1996). North American mammals data from this source yielded a test score of 0.159, a bit shy of 0.166. The latter guide, however, is substantially complete in taxa.

| Taxonomic group | K-S score | Taxonomic group | K-S score |
|---------------------------------|------------------|------------------------|------------------|
| Pteridophytes | 0.194 | Mammals | 0.130 |
| Angiosperms | 0.239 | Mammals | 0.122 |
| Arachnids | 0.272 | Pteridophytes | 0.322 |
| Herpetofauna | 0.110 | Angiosperms | 0.128 |
| Mammals | 0.287 | Arachnids | 0.032 |
| Pteridophytes | 0.270 | Herpetofauna | 0.108 |
| Angiosperms | 0.281 | Mammals | 0.157 |
| Arachnids | 0.123 | Mammals | 0.099 |
| Herpetofauna | 0.102 | Insects | 0.072 |
| Total score: 3.048 mean = 0.169 | | | |

Table 8.4. Contributions of continental-scale groups to K-S scores

The results of the full null hypothesis test appear in Table C2 (Appendix C), where the individual scores are organized on a taxon-within-taxon basis. Thus, “genera within families” means that for each taxonomic group with both genera and families represented, the corresponding test score is included under the given heading. Raw test scores that did not pass the Kolmogoroff-Smirnov test at a given level are underlined in the table.

In two of the tests, both involving extremely large numbers of taxa, the fits were exceptionally poor. In probing the reason for this, I discovered that the high score had less to do with a failure to match the shape of the proposed distribution, per se, and more to do with rather extreme variability from one abundance category to the next. By reducing all abundance numbers to ten percent of their original values, the effective sample size was reduced to lie within the upper end of the range of the other sample sizes.

| number predicted | 95% | 90% | 80% |
|-------------------------|-------|-------|-------|
| number passed | 92.7% | 87.3% | 81.8% |

Table 8.5. Comparison of K-S test scores with expected scores

New, much lower scores emerged, effectively dampening the variability. However, such data modification imposes a certain caution on conclusions influenced by either of the tests in question. Table 8.5 summarizes the outcome of the main experiment.

The agreement between the predicted and actual percentages is fairly good, with (percentage) differences of - 2.3 %, - 2.7 %, and + 1.8 %. Within each significance level some 55 tests were conducted. The passage of one more or one less score would, in each case, increase or decrease the percentage of passes by approximately 1.82%. To make the point of how close to optimal the foregoing results are, I have prepared another table to show what the results would be like if just one more test had passed in each of the first two score categories, along with one less passing in the third category.

| | | | |
|-------------------------|-------|-------|-------|
| number predicted | 95.0% | 90.0% | 80.0% |
| number passed | 94.5% | 89.1% | 80.0% |

Table 8.6. New scores with one test difference in each category

One is unlikely to get closer than the scores listed in Table 8.6 when 55 such tests are carried out.

8.4 Extinction and speciation in logistic-J communities

Extinction rates in experimental logistic-J communities such as the MSL system appear to be comparable with actual rates, although examples of the latter are not easy to find. As proxy for such real data, we must turn to experimental populations. As for speciation, there is a good reason to suspect that the multitude of small, somewhat isolated populations of a given species (all part of the same metapopulation) operate as a structural guarantee of that possibility, as will be seen.

8.4.1 Extinction rates in the MSL system

A great deal of research has been devoted to the problem of understanding extinction rates in natural communities, mainly through work with experimental populations. It is well understood that population size and “demographic stochasticity” are the dominant factors in determining the fates of individual populations. (Griffen and Drake, 2008)

Experiments with the Multi-species Logistic system have revealed that, depending on the value of ϵ or, equivalently, the minimum abundance, an isolated community will inevitably begin to lose species, owing to stochasticity within smaller populations. Yet the value of N will remain more or less constant. This can only mean, in this context, that the value of ϵ will continue to climb, while the minimum abundance (in the canonical sense) creeps slowly higher. As might be expected, extinctions become increasingly rare over time, as the community settles into a quasi-

equilibrium.

The actual extinction rates in the MSL system appear to be consistent with what we observe in real communities. Perhaps the most famous example of the latter was described by MacArthur & Wilson (1967) in their well-known study of island biogeography. When the island of Krakatoa exploded in 1883, all (visible) life on the island was presumably extirpated by lava and hot ash. But within a few decades some 30 species of birds had established themselves on the island, along with many plant species. At each visit, ecologists found that a few species were no longer present, but that new ones had arrived, more or less preserving the equilibrium number. MacArthur and Wilson estimated an average annual extinction rate of about 1.15 percent of species per year.

Accordingly, I set up the MSL system with 30 species having an average population of 200 individuals each (an estimate), then ran it to equilibrium (6000 cycles) at which point, of course, some bird populations were greater than 200 and others considerably smaller. At this point, the extinction switch was automatically turned on and, over the next 1000 cycles, the artificial community typically lost 9 species, a function of this particular combination of richness and abundance.

To discover roughly how much time in the life of a bird community 1000 cycles might represent, I used the following approximate reasoning: One cycle involves 100 iterations of the basic operation of trophism and reproduction. If 100 individuals reproduce in that time, it would represent all the females in a total population of 200 birds (breeding pairs). We may therefore estimate that each cycle was equivalent to turnover in a single species and that 30 cycles could therefore be taken as equivalent to one breeding cycle for all the birds on the island. Taking 30 cycles as equivalent to a single year, 1000 cycles would represent $1000/30 = 33.3$ years. Thus, in the avian equivalent of 33.3 years, the MSL system produced 9 extinctions, with the heaviest losses initially and declining losses thereafter. The initial loss rate was about 0.9 percent per year. If new species had been supplied to the system at this rate, the total number of species would of course remain approximately the same and the behavior of the system under these circumstances could hardly be distinguished from its behavior with the extinction switch simply turned off.

The average turnover of 1.13 percent on Krakatoa noted by MacArthur and Wilson is certainly in the same order of magnitude as the 0.9 percent produced by the MSL system under comparable circumstances. Indeed, given the relatively unsophisticated nature of the basic MSL system (not to mention the back-of-the-envelope approach taken here), the agreement is acceptably close. For example, since Krakatoa is well inside the tropical zone, one could allow two breeding cycles per year and arrive at a turnover rate of 1.8 percent, somewhat higher than the estimate of 1.13.

8.4.2 Speciation in stochastic communities

The fields of Ecology and Evolution are generally both represented in biology departments around the world. They often occur together as a single sub-unit. The reason for this is simple. Evolutionary theory must draw heavily on ecology in order to reconstruct the biota of past epochs. Limits to ecological knowledge automatically becomes limits to evolutionary knowledge.

The problem of speciation is a core topic within evolutionary biology. The study by Rice and Hostert (1993) summarizes research up to the date of that publication, where several theoretical models of speciation are evaluated in the light of breeding experiments with organisms ranging from vinegar flies to fish. It may be that, in spite of varying degrees of support offered by experimental evidence, all proposed mechanisms have a role to play, even if some play only a minor role. This situation offers an ironic contrast to the plethora of proposals for species abundance distributions where none may have a role to play. In any event, most of the proposed mechanisms depend on isolation of populations and acknowledge, implicitly or explicitly, that smaller populations are at some point necessary.

A large population can often be viewed as a metapopulation, with subpopulations distributed over a large area, perhaps even continental in scale. Logistic-J theory tells us that many of these subpopulations will be rather small but that, collectively, they may be large, qualifying them in a probabilistic sense for a speciation event. Of large populations that do not consist of isolated pockets, Mayr (1970) writes:

“The real problem of speciation is not how differences are produced but rather what enables populations to escape from the cohesion of the gene complex and establish their independent identity. No one will comprehend how formidable this problem is who does not understand the power of the cohesive forces in a coadapted gene pool”.

When a small local population receives a mutation, it has a good chance of spreading to all members of the population. Most often, the mutation is deleterious or makes no difference to the survival of the small population. Even if the mutation is advantageous, the population may well become extirpated. But if the population survives and ultimately grows, it carries the mutation with it. Eventually, the larger population replaces adjacent populations or remixes with them, introducing the new gene(s) into the larger pool and adding to the richness of the collective genome.

It may now be hypothesized that, over a long period of time, a species may vanish or become the host of a new species with approximately equal probability. If the metapopulation consists of a single pocket of individuals, extinction is a likely outcome. But if the metapopulation consists of many such pockets, speciation is a more likely outcome, making the low abundance end of the logistic-J distribution, not only the doorway to extinction but the very cradle of evolution.