Self-Organisation and Evolution of Social Systems

Edited by

CHARLOTTE K. HEMELRIJK

University of Groningen

CAMBRIDGE
UNIVERSITY PRESS

9

Natural selection and complex systems: a complex interaction

DAVID SLOAN WILSON
Binghamton University

In their book *Darwinism Evolving*, Depew and Weber (1995) develop the thesis that evolutionary theory has been reformulated several times to keep pace with advances in knowledge about the physical world. Darwin's Newtonian formulation was replaced by a probabilistic formulation early in the twentieth century. According to Depew and Weber, the new science of complexity will force yet another formulation, which is taking place during our time.

This general thesis may well be correct but Depew and Weber's specific account of the relationship between evolution and complexity leaves much to be desired (Wilson, 1995). They largely accept the polemic view of Gould and Lewontin (1979) at face value, arguing that natural selection is far more constrained and adaptations less common than claimed by proponents of the so-called adaptationist programme. Complexity is viewed as something that stamps its own properties on organisms and resists the modifying effects of natural selection.

Depew and Weber are not alone in this view. Many complexity theorists and writers seem to parade under the banner 'Darwin is dead! Long live complexity!' The following passage by Kauffman (1993: 24) provides one example:

> In short, if selection is operating on systems with strongly self-organized properties that are typical of the ensemble being explored, then those properties simultaneously are the proper null hypothesis concerning what we would expect to find in the absence of selection and may be good predictors of what we will observe even in the presence of continuing selection. In brief, if selection can only

Self-Organisation and Evolution of Social Systems, ed. Charlotte K. Hemelrijk.
Published by Cambridge University Press.

> slightly displace evolutionary systems from the generic properties of the underlying ensembles, those properties will be widespread in organisms not because of selection, but despite it.

This view of evolution and complexity suffers from being – well, too simple. More of one means less of the other. Complexity replaces natural selection as an explanatory principle rather than interacting with it. In this chapter I will present three case studies that illustrate a more complex and synergistic relationship between natural selection and complex systems. The synergy occurs because complex systems have profound effects on phenotypic variation and heritability, the two basic ingredients of natural selection.

Case study 1: radical epistasis and the genotype–phenotype–fitness relationship

Many problems in evolutionary biology are represented by the metaphor of adaptive landscapes, which invokes an image of peaks of high fitness separated by valleys of low fitness. Natural selection is envisioned as a hill-climbing process that moves populations to the tops of local peaks but is unable to move from low peaks to higher peaks that are separated by valleys. In addition, this metaphor makes it difficult to see how a sexually reproducing population can occupy more than one peak, because when individuals that occupy two peaks mate with each other, their progeny will be intermediate and will occupy valleys of low fitness (with the exception of complete dominance in single-locus models). This metaphor has influenced thinking on speciation (Otte and Endler, 1989), genetic polymorphisms (e.g. Charlesworth and Charlesworth, 1975), recombination (Bell, 1982; Michod and Levin, 1988), and genetic load (Wallace, 1991).

I call the adaptive landscape a metaphor because it has proven difficult to model rigorously. Provine (1986) showed that at least three incompatible versions exist, in which a point on the landscape represents (a) the mean fitness of the population, (b) the fitness of a single genotype, or (c) the fitness of a single phenotype. Even Sewall Wright, the primary inventor of the concept, used all three versions without realising that they cannot be easily related to each other. To make matters worse, a variety of factors such as density- and frequency-dependent selection can turn the landscape into a seascape with undulating peaks and valleys and can prevent natural selection from climbing hills, which surely is an essential part of the metaphor!

The passage by Kauffman quoted above is based in part on a model in which each genotype in a multi-locus system is assigned a phenotype, which in turn is assigned a fitness. Complex (epistatic) genetic interactions are modelled by making the genotype–phenotype relationship non-additive. The ultimate in epistasis is for a single allele substitution at any locus to create a new phenotype that is completely independent of the old phenotype. Fortunately, this kind of radical epistasis is easy to model, by simply defining a phenotypic parameter space and randomly assigning a value to each genotype in a multi-locus system. Kauffman showed that this kind of radical epistasis creates a 'rugged' fitness landscape in which the population can easily become trapped on peaks of low fitness and cannot find its way to higher peaks. Furthermore, a population cannot occupy more than one peak because mating and recombination create maladaptive combinations.

Kauffman's model complicates the genotype–phenotype relationship but keeps the phenotype–fitness relationship simple by assigning a constant fitness value to each phenotype. In contrast, Wilson and Wells (1994) complicate the phenotype–fitness relationship by making fitness density-dependent. The phenotypic parameter space consists of six values (5–10) that represent either a two-peak landscape (phenotypes 5&6 and 9&10 are peaks of high fitness separated by 7&8 as a valley of low fitness), or a three-peak landscape (odd-numbered peaks separated by even-numbered valleys). Valley phenotypes have a constant low fitness (= zero in most simulations) but the fitness of peak phenotypes declines with the number of individuals occupying the peak. Negative density- and frequency-dependent forces are extremely common in nature so these assumptions are at least as realistic as assuming a constant fitness for each phenotype.

Given these assumptions, the maximally adaptive genetic system would create phenotypes that occupy the peaks in equal proportions and miss the valleys. Of course, this is not possible if mating is random and the genotype–phenotype relationship is additive. In this case, phenotypes that occupy different peaks mate and produce intermediate offspring that fall into the valleys. However, something quite remarkable happens when the genotype–phenotype relationship is made maximally epistatic. In this case, a subset of genotypes evolves that satisfies two criteria: (a) all genotypes occupy adaptive peaks, and (b) all genotypes give rise to each other by recombination. Thanks to epistasis, the progeny of two individuals, which are genetically intermediate by definition, need not be phenotypically intermediate, which makes it possible for the genetic system as a whole to occupy more than one peak while missing the valleys.

Because a multi-locus genetic system contains many genotypes, it potentially contains many subsets that satisfy the two criteria outlined above. A single genetic system can adapt itself to both a two-peak and a three-peak adaptive landscape merely by selecting different subsets of genotypes. It is even possible to 'train' a radically epistatic genetic system to adapt to long-term evolutionary change in which the adaptive landscape switches periodically between a two-peak and a three-peak landscape. In this case, a subset of genotypes can evolve that satisfies the two criteria for one landscape, and then satisfies the criteria for the other landscape with a single mutational change (unpublished simulations).

To summarise, the same radically epistatic genotype–phenotype relationship that constrains natural selection in Kauffman's model endows natural selection in our model with a remarkable ability to mould itself to multi-peak adaptive landscapes and swiftly adapt to long-term recurrent changes in the environment. The reason is obvious, if only in retrospect: a radically epistatic system allows almost any combination of phenotypes to be produced by a randomly mating subset of genotypes, without the production of intermediate forms that has always made evolution in multi-peak adaptive landscapes seem problematic.

Case study 2: group selection in humans and other animals

Natural selection within a single group is insensitive to the welfare of the group (Sober and Wilson, 1998). A solid citizen who behaves 'for the good of the group' will not increase its proportional representation within the group, and will even decline in frequency (ultimately going extinct) if providing the public good requires a private cost in time, energy or risk. Freeloaders are equally or more fit than solid citizens within groups. On the other hand, groups of solid citizens are more fit than groups of freeloaders. If there are many groups that vary in their proportions of solid citizens and freeloaders, the solid citizens can evolve by group-level selection despite their selective disadvantage within groups.

However, why should groups vary in their proportions of solid citizens and freeloaders? In the past, models that attempt to answer this question have been based on simple genotype–phenotype relationships, such as a single-locus model in which one allele codes for the solid citizen and the other codes for the freeloader. In this case, genetic and phenotypic variation become tightly linked to each other and the partitioning of variation within and among groups is determined straightforwardly by sampling error, declining precipitously with the number of individuals that independently colonise the groups. Since group selection requires variation among groups, the implication is that the importance of group selection declines rapidly with initial group size. These models contributed to the widespread rejection of group selection as an important evolutionary force in the 1960s. Few people questioned whether the verdict might have been based on the assumption of a simple genotype–phenotype relationship.

Subsequent research over the last few decades has amply shown that the rejection of group selection was premature. In the first place, real groups simply do not conform to expectations based on sampling error; they often vary greatly in their phenotypic properties, even when initiated by many individuals. In the second place, selection experiments reveal that phenotypic variation among groups is usually heritable, sometimes even more heritable than phenotypic variation among individuals within groups (Goodnight and Stevens, 1997). Group selection happens, in both nature and the laboratory, so why did the initial models fail so miserably to represent the real world?

The answer can be traced to the assumption of a simple genotype–phenotype relationship. More complex relationships can cause random genetic variation among groups to give rise to highly non-random phenotypic variation. There are many specific ways to do this but all of them can be understood as forms of sensitive dependence on initial conditions. A population of groups can be regarded as replicate biological systems that initially vary, if only at random. If the biological systems are complex in their genetic and social interactions, then the initial differences will not stay small but will provide the basis for larger differences – the 'butterfly effect' that is so well known for complex physical systems. Replicate complex systems are *expected* to vary, no matter how small their initial differences. We might wonder if the differences that develop are heritable, but we shouldn't be surprised by the raw fact of phenotypic variation among groups whose elements interact in a complex fashion, even if the initial variation is random.

I will provide one example of the decoupling of genetic and phenotypic variation, although many could be chosen. Wilson and Kniffin (1999) present a multilevel selection model in which genes code for social transmission rules rather than directly for behaviours. One rule is for conformity and can be stated formally as: 'converge upon a single behaviour, chosen randomly from the initial behaviours in the group'. If everyone in a group follows this rule, and if they enter the group exhibiting different behaviours, then all of them adopt a single behaviour chosen at random from the initial behaviours. There are a variety of specific mechanisms that might accomplish this but we need only focus on the final effect of behavioural uniformity within the group.

Now let's expand our scope to consider a large population of individuals who follow this rule. A fraction p behave like solid citizens while the remaining $(1 - p)$ behave like freeloaders. They form into many groups of size n at random. At first behavioural variation among groups follows the binomial distribution, but then the social transmission rule takes over and all groups become behaviourally uniform. Notice that the value of p has not changed in the total population because each group converges on one or the other behaviour at random, but the partitioning of the variation has changed dramatically. Initially most of the variation was within groups, whereas all of it becomes variation among groups thanks to the transmission rule. Of course, the new partitioning of variation is maximally favourable for group selection. The groups of solid citizens do well, the groups of freeloaders do poorly, there is no exploitation within groups, and so the differential fitness of groups increases the frequency of solid citizens (*p*) in the global population. If the groups break up and the individuals (or their progeny) form new groups at random, the transmission rule again eliminates variation within groups and maximises variation among groups, and the cycle repeats itself.

So far, this model shows how a population can be genetically uniform but phenotypically variable, and how large phenotypic differences among groups can develop even when the groups are randomly formed. However, it remains to show how this particular genetically encoded transmission rule can evolve in competition with other transmission rules. Wilson and Kniffin (1999) consider three rules in addition to the one already described: (a) a rule that resists changing its behaviour after entering a new group, (b) a rule that is biased toward becoming a solid citizen, and (c) a rule that is biased toward becoming a freeloader. A group formed by individuals with different transmission rules will not necessarily converge upon a single behaviour. Thus, there will be mixed groups in which freeloaders profit at the expense of the solid citizens. Nevertheless, the phenotypic variation among groups can still be non-random, despite the fact that the groups are randomly formed (Fig. 9.1). Moreover, the relationship between phenotypes and genotypes has become complex; a given genetic type can become either behavioural type. Even when the model includes 'recalcitrant' genotypes biased toward freeloading, this complex phenotype–genotype relationship can nevertheless favour the evolution of transmission rules that are not only conformist, but biased toward solid citizen behaviours. In more familiar terms, evolution in randomly formed groups can evolve the psychological disposition to figure out what's good for the group and make sure that everyone does it. More generally, this model is only one of many in which complex genotype–phenotype relationships make group selection a potent evolutionary force, even in groups initiated at random by large numbers of individuals.

Case study 3: ecosystem selection

The idea that entire ecosystems can evolve into adaptive units has been regarded as among the most extravagant claims of holistically minded biologists. Nevertheless, the same principles that I outlined in case study 2, which enable complex single-species systems to evolve into adaptive units, also apply to ecosystems consisting of multiple species interacting with each other and their abiotic environment.

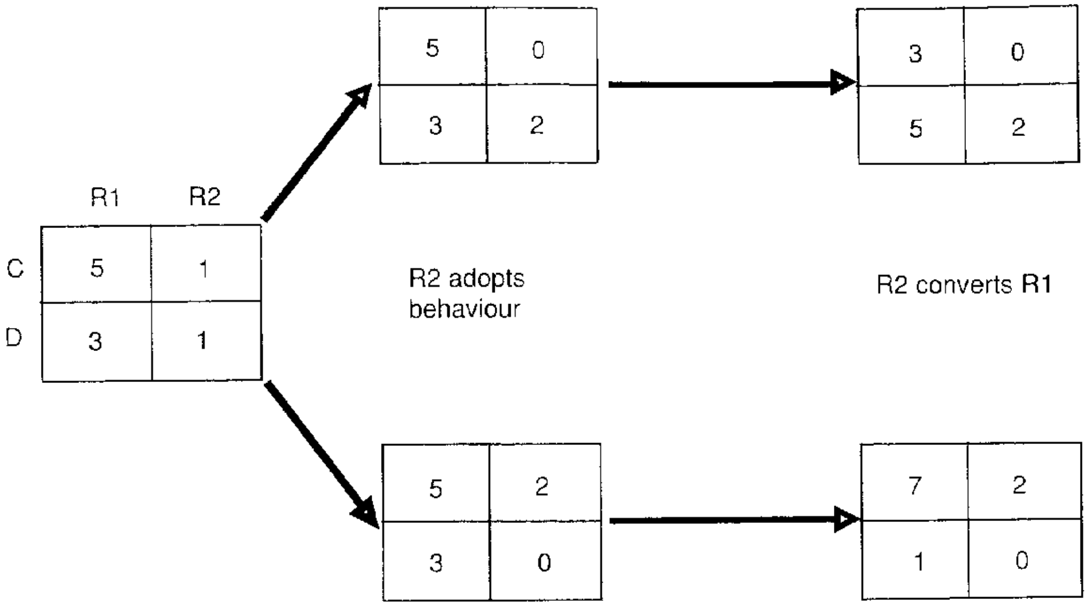

Figure 9.1 An interaction between behaviours (C and D) and transmission rules (R1 and R2) increases variation among groups. The transmission rules are assumed to be genetically encoded while the behaviours are acquired on the basis of the transmission rules. C and D are the behaviours 'cooperate' and 'defect' in a standard 'prisoner's dilemma' model. Individuals with transmission rule R1 retain their behaviour unless influenced by another transmission rule. Individuals with transmission rule R2 converge upon a single behaviour chosen at random from the initial behavioural composition of the group. In addition, each R2 individual is able to 'convert' a single R1 individual. The group starts with the initial composition of behaviours and transmission rules shown in the left-hand box and follows one of the two trajectories shown by the arrows, depending upon whether the R2s converge upon C or D. In a global population with many groups, the effect of these interactions is to decrease behavioural variation within groups and increase behavioural variation among groups, favouring the evolution of cooperation. From Wilson and Kniffin (1999).

Of course, ecosystems are vastly more complex than any single species. This is no argument against ecosystem selection, however, since the basic point of this essay is to show that complexity can *facilitate* the natural selection process. Consider the community of microbes, fungi and invertebrates that inhabit the root zone of plants. Each plant and its associated fauna can be regarded as a local ecosystem, whose members interact primarily among themselves and with the physical environment in the immediate vicinity of the plant. These local ecosystems will surely vary in their species composition and in the genetic composition of the component species. Even if the initial variation among ecosystems is slight, sensitive dependence is likely to magnify these differences, as described for single-species groups in case study 2. The fact that local ecosystems are more complex than local groups of a single species makes them even more likely to diverge on the basis of initial differences. Variation in the species/genetic composition of the local ecosystems is likely to have an effect on important ecosystem processes such as nutrient cycling, oxygen availability, modification of the physical substrate, and so on. Solid citizens and freeloaders are at least as relevant to ecosystem processes as to the dynamics of single-species groups. Local ecosystems that function well, perhaps by helping 'their' plant to grow vigorously and provide them with abundant resources, could then contribute differentially to the formation of future local ecosystems.

Simulation models (Wilson, 1992; Johnson and Boerlijst, 2002; Johnson and Seinen, 2002), laboratory experiments (Swenson *et al.*, 2000a, b), and even a limited amount of field research (Wilson and Knollenberg, 1987) suggest that this scenario is fully plausible, no matter how far-fetched the concept of ecosystem selection may have appeared in the past. One key insight is that most ecological processes are local by their nature, which makes it possible to envision a large-scale ecosystem such as a forest as a population of many small local ecosystems. Once a population of ecosystems becomes plausible, complex interactions make variation among ecosystems almost inevitable. The ecosystems need not be separated by discrete boundaries. Simulation models show that even when species are placed on a uniform grid, such that each individual is at the centre of its neighbourhood, complex localized interactions result in spatial heterogeneity. If variation among patches on a landscape such as this has functional consequences, the stage is set for the successful patches to persist and spread at the expense of less successful patches.

The simulation model of Wilson (1992) illustrates how complex interactions can produce heritable variation at the level of multi-species communites. Local communities are colonised at random from a global pool of 10 species and interact for a period of time before dispersing back into the global pool. Species interactions are modelled with Lotka–Volterra equations and a three-dimensional matrix of interaction terms randomly and uniformly distributed between 0 and 2, where $a_{i,j,k}$ is the effect of species i on species j in the presence of species k. For every iteration of the difference equations, the three-dimensional matrix is collapsed into a two-dimensional matrix by averaging the 10 $a_{i,j}$ terms, weighted by the relative proportions of the 10 species in the patch. In this fashion, the pairwise species interactions depend upon the background composition of the local community.

These complex interactions cause initially random variation in the species composition of local communities to become highly non-random, as shown in Fig. 9.2. The upper graph shows the relative proportion of the 10 species in the global dispersal pool over the course of 21 cycles of local community

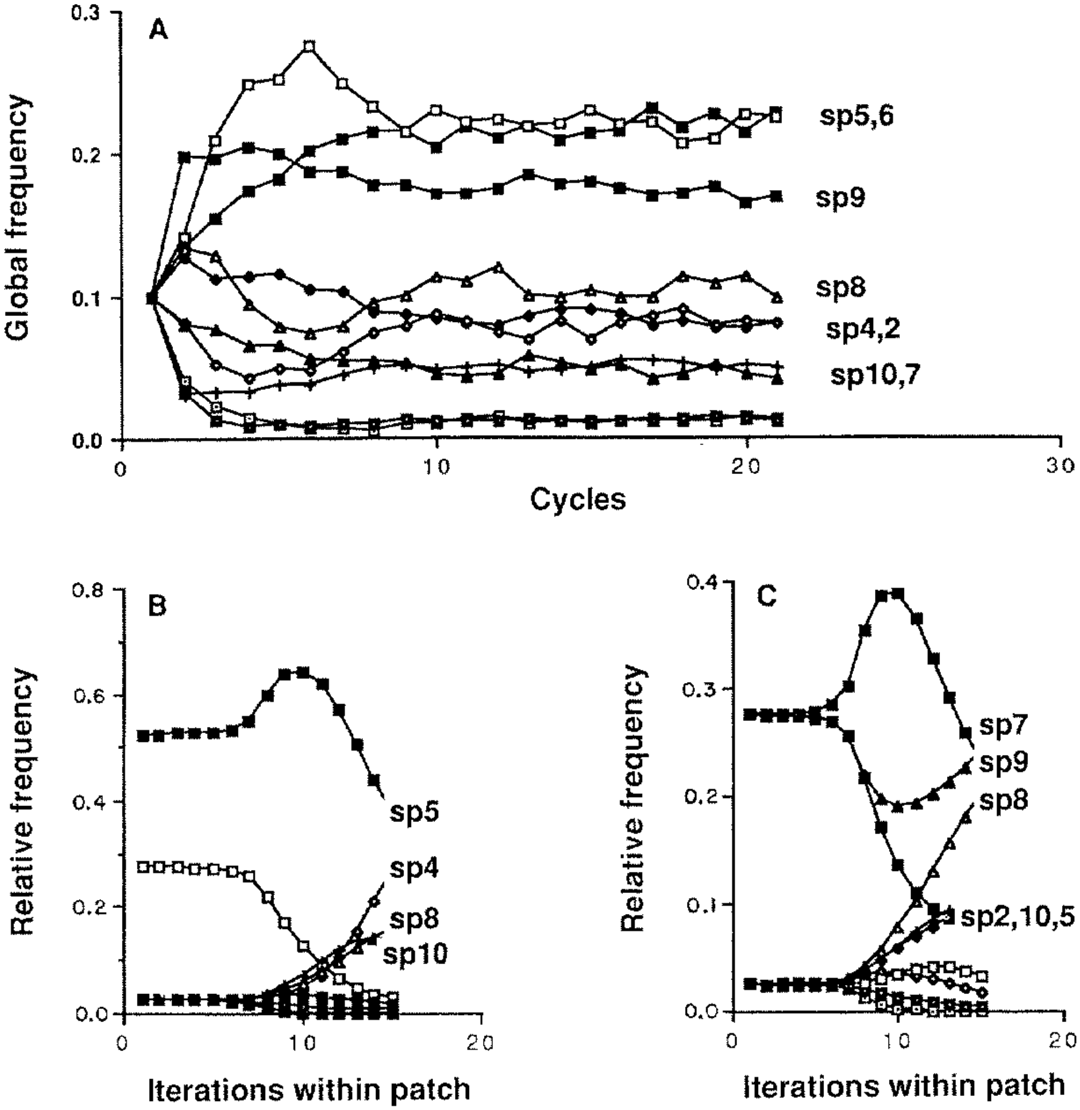

Figure 9.2 Community-level selection in a simulation model, in which local communities (patches) are colonised at random from a global pool of dispersers, and the species within each patch interact for a period of time before dispersing back into the global pool. The simulation is initiated with all species in equal frequency. (A) Changes in frequency in the global pool over 21 cycles of patch formation and dispersal; (B) and (C) the dynamics of two local patches. From Wilson (1992).

formation and dispersal. Community dynamics at the global scale looks fairly orderly and is governed by negative frequency-dependent selection. In other words, if any given species is displaced from its characteristic frequency, it will return to that frequency. The two lower graphs show the dynamics of a sample of local communities drawn at random from the global pool after it has reached its (rough) equilibrium. Community dynamics at the local scale is highly non-equilibrium, with the species composition embarking on different trajectories depending upon initial conditions. Nevertheless, the orderly dynamics at the global scale is produced by the combined non-equilibrium dynamics at the local scale.

The result of all this is that many species are maintained in the global pool and that local communities perpetually vary in their species composition. Now what happens if the local communities differentially go extinct and fail to contribute to the global pool (community-level selection)? For example, what happens when local communities are made to go extinct in direct proportion to the abundance of species 5 in Fig. 9.2? Not only does species 5 become less abundant in the global community, but the other species change their relative proportions as well, as shown in Fig. 9.3. The reason is that species 5 was abundant in some local communities and rare in others due to the presence of other species that supported or inhibited the growth of species 5. Community-level selection not only differentially removed species 5 but also its supporting cast (including species 4 and 19), while favouring the species that prevent species 5 from becoming abundant in the first place (including species 9, 2 and 7). Before community-level selection, we had a species that drove the entire local community extinct. After community-level selection, we have local communities that drive the offending species extinct or at least to low levels of abundance.

Empirically, the laboratory experiments of Swenson *et al.* (2000a, b) impressively demonstrate how the butterfly effect can operate in biological systems. Microcosms that were made as physically identical as possible were inoculated with millions of microbes from a single well-mixed source. Based on sampling error, there was negligible initial variation in the species composition of each microcosm, yet they became very different over a period of only a few days (which equals many microbial generations) and the differences had important consequences for ecosystem properties such plant growth, pH of the environment or degradation of a toxic compound (depending on the experiment). When these properties were used as a basis for selecting whole microcosms to colonise a new set of microcosms, there was a response to selection, demonstrating that the kind of phenotypic variation created by complex interactions can be heritable (Figs. 9.4 and 9.5).

One way to think about this experiment in terms of complex systems theory is to imagine a parameter space with $(S + 1)$ dimensions, where each axis represents the proportion of one species in an S-species ecosystem and the $(S + 1)$th axis measures the ecosystem property that serves as the basis for selection (e.g. pH of the medium). At the beginning of the experiment described above, the microcosms are represented by a small cloud of points in the parameter space, because they are nearly identical in their species composition except for

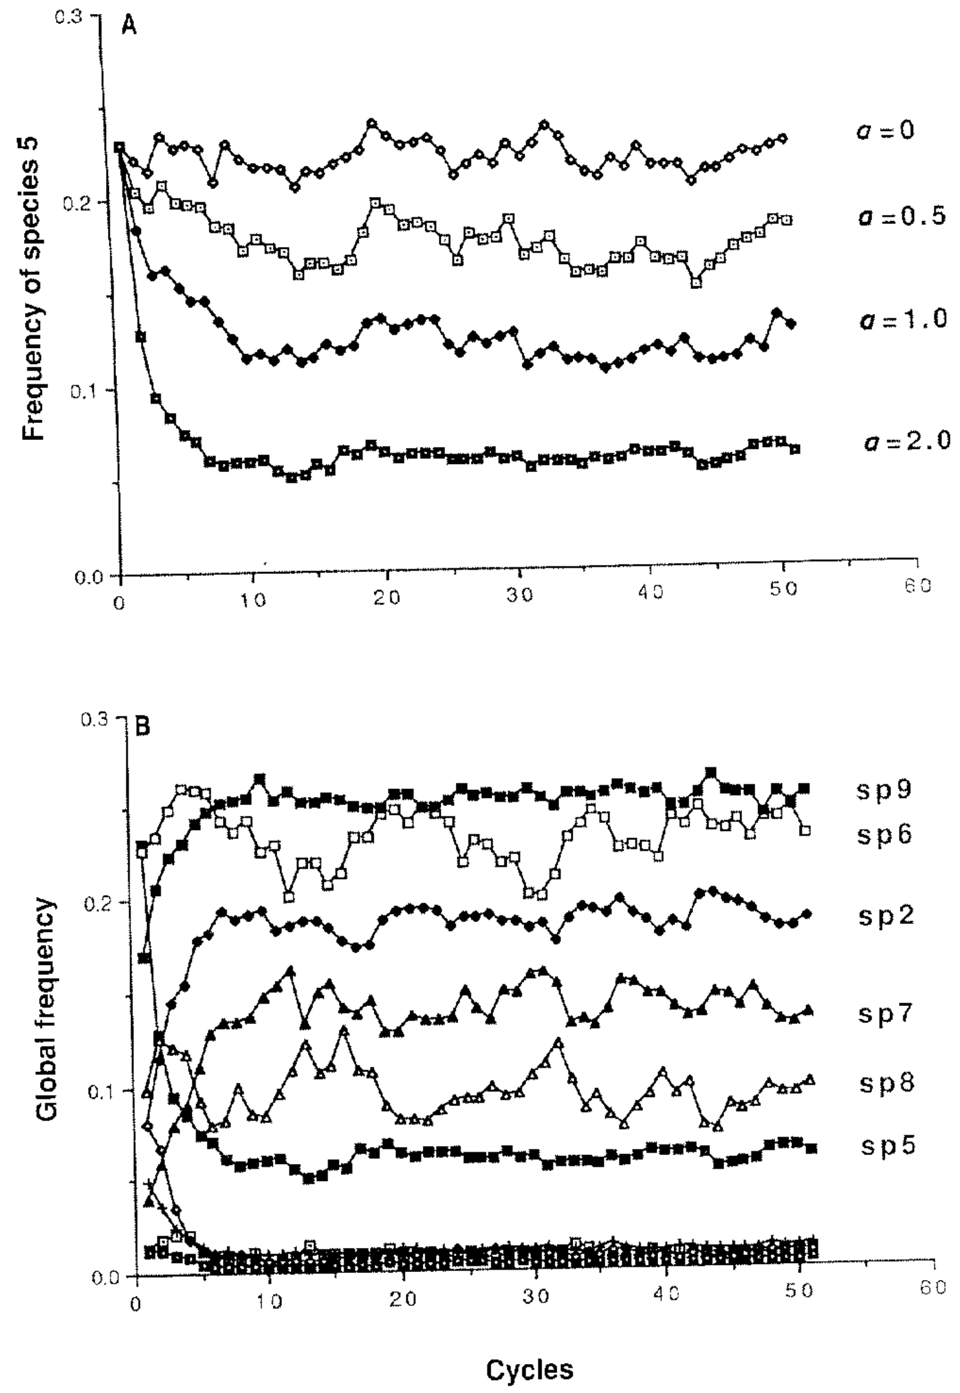

Figure 9.3 A continuation of Fig. 9.2, with the added factor that local communities go extinct in direct proportion to the abundance of species 5 just prior to dispersal. Note that species 5 is one of the most abundant species in the global community in Fig. 9.2, although its abundance varies at the patch level: (A) shows the change in the global frequency of species 5 as the intensity of extinction at the group level (governed by the a term) increases; (B) shows the frequency of all 10 species when $a = 2$, which can be compared to Fig. 9.2 when $a = 0$. Community-level selection changes the composition of the entire community, reducing not only species 5 but also its 'supporting cast'. See text and Wilson (1992) for more details.

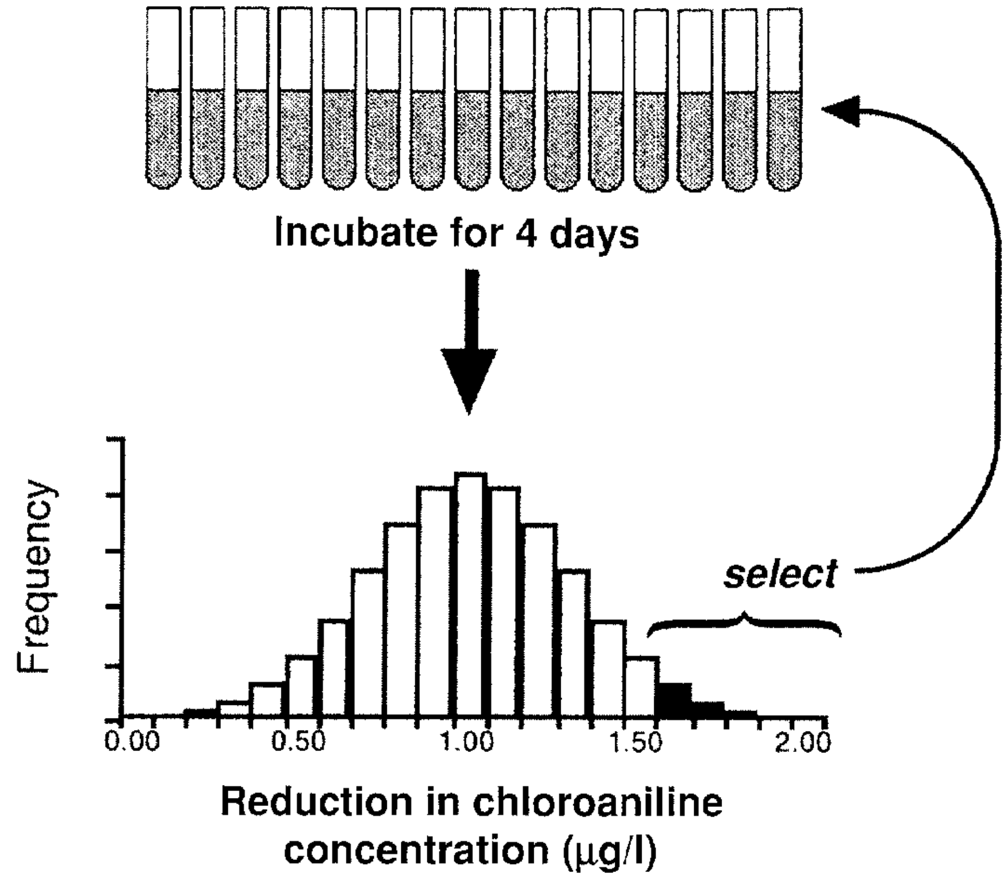

Figure 9.4 Schematic diagram of an ecosystem-level selection experiment. Test tubes with microbial growth medium and a toxic compound added (chloroaniline) are inoculated with a naturally occurring microbial community from a single source. Sensitive dependence causes the microcosms to diverge in their species composition, with corresponding differences in chloroaniline degradation. The microcosms that most effectively degrade chloroaniline are used as 'parents' to create a new 'generation' of microcosms. Note that the microcosms are being selected on the basis of a property of the abiotic environment (chloroaniline concentration) that is influenced by biotic processes, which is why this experiment counts as an example of ecosystem selection. From Swenson *et al.* (2000a).

sampling error. As complex interactions take place within each microcosm, sensitive dependence causes them to embark upon separate trajectories and the cloud of points expands. To the degree that species composition influences the ecosystem property that serves as the basis for selection, variation will increase along the $(S + 1)$th axis as well. At the end of the first generation, the microcosms at one end of the distribution on the $(S + 1)$th axis are selected and used as the 'parents' which are mixed together and used to inoculate a second generation of microcosms. Variation at the beginning of the second generation is again represented by a small cloud of points, but the location of this cloud is

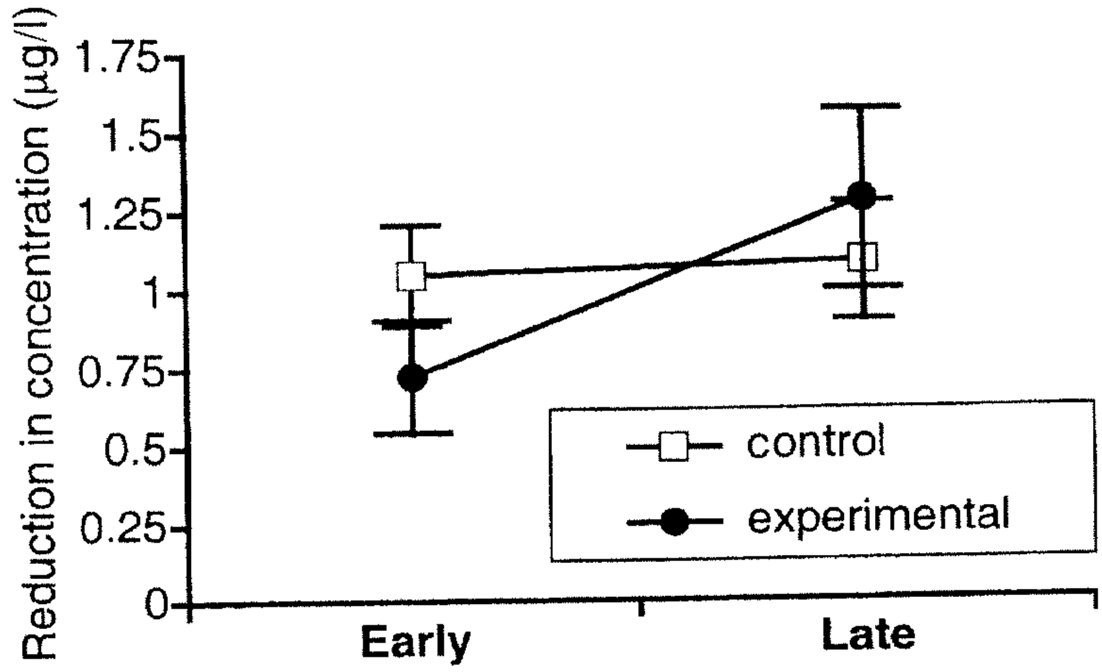

Figure 9.5 Response to selection in an ecosystem selection experiment. Four experimental lines selected for chloroaniline degradation increased their degradation ability on average, compared to four control lines in which the microcosms were selected at random. The average degradation ability of the experimental lines was (non-significantly) lower than the control lines at the beginning of the experiment by chance. From Swenson *et al.* (2000a).

different from that for the first generation. In this fashion, the experiment can be envisioned as a way of searching a very large parameter space for regions that create the phenotype being selected (e.g. high pH), similar to computer algorithms that use multiple agents to search large parameter spaces. However, regions of the parameter space that create the selected phenotype must also be locally stable for there to be a response to selection. In the absence of local stability, the ecosystems that were selected during one generation will give rise to ecosystems during the next generation that embark on trajectories of their own, with different effects on the ecosystem property that serves as the basis of selection. To summarise, if the parameter space includes regions that (a) create the phenotype being selected and (b) are locally stable, then the process of ecosystem selection can 'find' these regions. Switching back from the language of complex systems theory to the language of evolutionary theory, there will be phenotypic variation at the ecosystem level and some of it will be heritable. Notice that ecosystem selection can take place purely by changing species composition without any genetic changes in the component species. The local ecosystem has truly become the analogue of an organism and species have become the analogue of genes. However, the response to ecosystem selection can include genetic changes within species in addition to changes in species composition.

Conclusion

Scientific inquiry often begins at the simplest possible level and adds complexity only when needed. Thus it is not surprising that evolutionary and ecological models began by assuming simple interactions. However, the inadequacies of these assumptions are becoming apparent and further progress will require a greater appreciation of complexity in various forms. In this sense, Depew and Weber (1995) are right that a new formulation of evolutionary (and ecological) theory is in order. However, the new formulation will not merely assign a smaller role to natural selection and a larger role to the properties of complex systems that resist selection. Some resistant properties undoubtedly exist, but more important will be the properties of complex systems that facilitate natural selection. In each of my case studies, complex systems provided far more promising raw material for the sculpting action of natural selection than simple systems. How fitting that the relationship between natural selection and complexity is proving to be complex.

References

Bell, G. (1982). *The Masterpiece of Nature*. Berkeley, CA: University of California Press.

Charlesworth, D. and Charlesworth, B. (1975). Theoretical genetics of Batesian mimicry. I. Single-locus models. *J. Theoret. Biol.* **55**, 283–303.

Depew, D. J. and Weber, B. H. (1995). *Darwinism Evolving: Systems Dynamics and the Genealogy of Natural Selection*. Cambridge, MA: MIT Press.

Goodnight, C. J. and Stevens, L. (1997). Experimental studies of group selection: what do they tell us about group selection in nature? *Am. Naturalist* **150**, S59–S79.

Gould, S. J. and Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist program. *Proc. Roy. Soc. London* B **205**, 581–598.

Johnson, C. R. and Boerlijst, M. C. (2002). Selection at the level of community: the importance of spatial structure. *Trends Ecol. Evol.* **17**, 83–90.

Johnson, C. R. and Seinen, I. (2002). Selection for restraint in competitive ability in spatial competition systems. *Proc. Roy. Soc. London B* **269**, 655–663.

Kauffman, S. A. (1993). *The Origins of Order*. Oxford: Oxford University Press.

Michod, R. E. and Levin, B. R. (eds.) (1988). *The Evolution of Sex*. Sunderland, MA: Sinauer Associates.

Otte, D. and Endler, J. A. (eds.) (1989). *Speciation and its Consequences*. Sunderland, MA: Sinauer Associates.

Provine, W. B. (1986). *Sewall Wright and Evolutionary Biology*. Chicago, IL: University of Chicago Press.

Sober, E. and Wilson, D. S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.

Swenson, W., Arendt, J. and Wilson, D. S. (2000a). Artificial selection of microbial ecosystems for 3-chloroaniline biodegradation. *Envir. Microbiol.* **2**, 564–571.

Swenson, W., Wilson, D. S. and Elias, R. (2000b). Artificial ecosystem selection. *Proc. Natl Acad. Sci. USA* **97**, 9110–9114.

Wallace, B. (1991). *Fifty Years of Genetic Load: An Odyssey*. Ithaca, NY: Cornell University Press.

Wilson, D. S. (1992). The effect of complex interactions on variation between units of a metacommunity, with implications for biodiversity and higher levels of selection. *Ecology* **73**, 1984–2000.

(1995). Book review: *Darwinism Evolving: Systems Dynamics and the Genealogy of Natural Selection*. *Artif. Life* **2**, 261–267.

Wilson, D. S. and Kniffin, K. M. (1999). Multilevel selection and the social transmission of behavior. *Hum. Nature* **10**, 291–310.

Wilson, D. S. and Knollenberg, W. G. (1987). Adaptive indirect effects: the fitness of burying beetles with and without their phoretic mites. *Evol. Ecol.* **1**, 139–159.

Wilson, D. S. and Wells, A. (1994). Radical epistasis and the genotype–phenotype relationship. *Artif. Life* **2**, 117–128.