Last Will and Testament: Stephen Jay Gould's *The Structure of Evolutionary Theory**

Kim Sterelny[†]

I outline Gould's conception of evolutionary theory and his ways of contrasting it with contemporary Darwinism; a contemporary Darwinism that focuses on the natural selection of individual organisms. Gould argues for a hierarchical conception of the living world and of the evolutionary processes that have built that living world: organisms are built from smaller components (genes, cells) and are themselves components of groups, populations, species, lineages. Selection, drift and constraint are important to all of these levels of biological organization, not just that of individual organisms. Moreover, both drift and constraint are more important than orthodoxy supposes. While having some sympathy for both of these lines of argument, I argue that they are more problematic than Gould supposes, and that he understates the power and the heterogeneity of orthodox conceptions of life's evolution.

Stephen Jay Gould, **The Structure of Evolutionary Theory.** Cambridge, MA: Harvard University Press, (2002), xxii + 1433.

1. Darwin's Edifice. This massive book develops Gould's distinctive picture of evolutionary biology as it has been, as it is and as it should be. Gould's strategy is to identify the logical core of Darwin's own doctrine and chart its history. He takes that core to be a three-part thesis about natural selection. The first part is about *agency.* Almost always, natural selection acts on individual organisms rather than groups or species. The second part is about the *efficiency* of natural selection. Natural selection, working gradually and cumulatively, has the power to build novel biological or-

---

*Received November 2002.

[†]To contact the author, write to: Philosophy Program, Victoria University of Wellington, PO Box 600, Wellington, New Zealand, or Philosophy Program, RSSS, Australian National University, 0200, Canberra, ACT, Australia; e-mail: Kim.Sterelny@vuw.ac.nz or kimbo@coombs.anu.edu.au

Philosophy of Science, 70 (April 2003) pp. 255–263. 0031-8248/2003/7002-0002$10.00

ganizations. The third part is about selection's *scope.* Natural selection acts on individuals within a local population, and if suitable variation is present, selection can then alter the average phenotype of that population. These microevolutionary changes accumulate into the vast patterns in time and space revealed by paleobiology. For Darwin, these patterns are just summed, extrapolated microevolutionary events and nothing more.

Having identified the logical core of Darwin's own vision, in the first half of this book Gould charts the fate of these ideas up to the recent past. In doing so, he identifies enduring challenges to these three central themes. These challenges are persistent because, as Gould sees it, Darwin's edifice needs substantial rebuilding, though Gould insists (rightly) that his version of evolutionary theory remains recognizably Darwinian. I will begin by briefly outlining a rather caricatured contemporary version of Darwin's edifice before discussing two revisions. This textbook Neo-Darwinism interprets evolutionary change in the following way: (i) Organisms in a local population vary somewhat. (ii) Some of this phenotypic variation has fitness consequences: these explain differences in reproductive success in these populations. (iii) Some phenotypic differences are heritable. Offspring resemble their parents more than they resemble a randomly chosen member of their population, and in ways that are relevant to fitness. Local populations vary not just in their phenotypes but in their genotypes, and those genetic differences have systematic, though not simple, effects on the phenotypes of the organisms carrying those genes. (iv) Since in reproduction, organisms transmit some of their genes to their descendants, biased reproduction in the adult generation leads to biased gene transmission, and hence to a changed average phenotype in the offspring generation. For example, if there is selection for (large) size, and if there is variation in the population for genes that have systematic effects on size, the proportion of such genes will increase in the descendent population. (v) Evolutionary change is the iteration of this process. Changes in species and in lineages of species result from nothing but such population-level processes summed over space and time. In the rest of this notice, I shall discuss two important elements of Gould's reconstruction of this picture: the relationship between macro- and microevolution and the role of developmental biology. Constraints of space force me to ignore much that Gould discusses; most importantly, the interaction between levels of biological individuality.

2. Hierarchy and Extrapolation. The basic principles of evolution are neutral on the nature of individuals, lineages, and populations. So long as we have a population of well-defined individuals that produce others like themselves with differential success, then the population will evolve. The character of the individuals will change over time. Furthermore, organisms are both composed of individuals—cells and genes—and, arguably, are components of further individuals: groups, species and perhaps others. It is now uncontroversial that selection can act below the level of the organism, for there are genes that can increase their representation in the next generation by biasing the gene transmission process. For example, there are sex-linked genes that bias sex ratios. One of the most central themes of *Structure* is that organism collectives are important units of selection too. Gould endorses the recent revival of group selection, but his main interest is defending the importance of species selection. For while he concedes that there are as yet few well-confirmed empirical studies of species selection in action (709–710), and while he accepts that species selection plays little role in the evolution of the complex adaptive systems of individual organisms, he thinks that it has played a major role in generating evolutionary trends and other macroevolutionary patterns (731–735, 738–744).

In developing a hierarchical view of evolution, Gould's first important contribution to evolutionary biology, the hypothesis of punctuated equilibrium, plays an important role. According to that hypothesis, the typical life history of a species consists in the geologically rapid evolution of the phenotypic differences that distinguish it from its ancestor, followed by phenotypic stasis until the species disappears by extinction or by splitting into daughters. The extent to which this does describe the typical life history of a species is still open to debate. But if Gould is right, it sets up a challenge to the orthodoxy above. First, it underwrites the biological reality of species, showing how they have determinate boundaries in time (603–608, 714–720). Second, it directly challenges extrapolationism. For while Gould's case for the importance of species selection remains rather fragile, the case against extrapolationism is much stronger. The punctuated equilibrium pattern, if typical, strongly suggests that evolution at the species level is not a simple projection of evolution in local populations.

The problem is stasis. As Gould notes, the background extrapolationist expectation is one of gradual change, rather than long periods without change interrupted by geologically sudden bursts of species formation. Of course, there are plenty of microevolutionary processes that would result in stasis in particular cases. A species phenotype might be held fixed by stabilizing selection. Alternatively, as conditions change, a species might track its preferred habitat. Adaptation might be behavioral or physiological rather than morphological, and hence invisible to paleobiology. Even so, we would not expect stasis to be the *dominant pattern* of the fossil record, if that record is nothing but population-level evolution summed over large quantities of space and time. Microevolution takes place in populations, yet stasis is the result of properties of species, not populations. Gould argues (I think rightly) that stasis is likely to be the result of the division of species into populations, each of which is subject to somewhat different selection pressures. Adaptive change does take place in these local populations, but it is ephemeral, both because local conditions fluctuate, and because local populations have fluid boundaries, and migrations from outside dilute local adaptation. The structure of species acts as a brake. In most circumstances, local adaptation is neither entrenched nor accumulated (778–780, 798–801).

Gould is probably right in rejecting the strongest versions of extrapolationism. What though of species selection? Here I think his case is more fragile, on both conceptual and empirical grounds. If punctuated equilibrium describes the predominant pattern of a species' life history, species are indeed identifiable biological individuals, though whether they form populations is less clear. I shall focus not on that problem but on the notions of fitness and selection. On the face of it, there seems no difficulty in defining a notion of fitness appropriate to a species itself. A species' fitness is determined by its propensity to resist extinction and to give rise to daughters by budding or splitting. The fitness of species seems to be directly parallel to the fitness of asexual organisms. However, it turns out that identifying species fitness is not straightforward.

2.1. Drift versus Selection. The success and failure of individual organisms is not always the result of selection. Some mortality is the result of unhappy accident rather than ill design. Some organisms are fecund by luck, despite their design flaws. Hence *realized fitness* differs from *expected fitness.* The smaller the population, the more likely it is that the population will be affected by biological accidents: lucky survivors, unlucky deaths. Such unselected evolutionary changes are one aspect of drift. Gould expects drift to be important at the species level, for the number of species in a clade is small by comparison to the number of organisms in a species. It is more likely that the future of a lineage is affected by the unlucky extinction of a species, than it is that the future of a species is affected by the unlucky death of an individual organism (735–738).

So far so good. However, species population sizes are so small, and cycles of selection are so few, that we are in danger of losing the distinction between selective and accidental mortality. Richard Lewontin has pointed out the deep problems of this distinction. If we consider an organism's life history in all its contingent detail, everything that happens to it is predictable given its environment. So we lose the distinction between expected and realized fitness, and "the survival of the fittest" really becomes a tautology. Thus we abstract away from some of this detail. Only patterns in the experience of the organisms in the population count. The death of a Major Mitchell cockatoo from lightning strike is an unlucky accident, not selection, for that death is not systematic. That death stands in no regular relationship to the habits of the species. If, in contrast, some satin bowerbirds developed a preference for decorating their bowers with copper wire, thus making them more vulnerable to lightning strike, the frying of a particular bird would be selection. But systematic patterns in causes of mortality exist, it seems, only against a background of numbers and cycles of selection.

The problem, then, is that we have no principled way of distinguishing between a species' expected and actual fitness. Which details of a species' habitat and life history can we legitimately abstract away from, in determining a species' expected fitness? Was the extinction of New Zealand's moa species drift, with the arrival of the Maori the equivalent of the unlucky bolt of lighting? Would it have been drift, had there been only one species of moa, just as there was only one species of elephant bird on Madagascar, when humans arrived? This issue is of particular importance to Gould, for he emphasizes the role of mass extinction in shaping the tree of life. Gould has argued that those extinctions are neither genuinely random nor biased in favor of fitter, better-adapted species (a "fair game"). Rather, these are periods of different rules of failure and survival (1312–1320). Yet these distinctions require a solution to Lewontin's puzzle. No extinction is random, if every feature of a species' environment plays a role in determining its adaptive landscape. The distinction between "fair game" and "wanton" extinction collapses too. Perhaps the geographic and temporal location of a species—the location exposing it to a mass extinction event—is part of its selective environment. After all, the variability of the arctic climate is a salient part of the selective environment of a polar bear that must survive both the winter and the summer.

2.2. Selection and Sorting. Suppose we can solve this problem. Another lurks: the distinction between species sorting and species selection. Suppose we notice a difference between two lineages of birds. One appears successful. Over time, its geographical zone expands, it occupies more niche space, and both species diversity and absolute numbers go up. At the same time and place, another lineage is in decline. Its range shrinks and fragments. Its diversity falls. Even where they are still found, the animals themselves become less abundant. Should we take these contrasting patterns to show species-level selection in favor of one lineage, and against the other?

Gould notes that any such conclusion would be grossly premature. This pattern in species failure/success may well be a macroevolutionary effect of individual adaptation and maladaptation. Suppose the story is as follows. The dwindling lineage consists of birds that nest in tree hollows. And safe hollows are in increasingly short supply, because of the evolution and spread of a lineage of egg-eating snakes. In contrast, the successful lineage lays its eggs in woven cups which the birds suspend from the ends of twigs. These are too fragile to support a snake's weight, and hence the birds are safe from snake predation.

All defenders of a hierarchically expanded theory of evolution recognize Vrba's distinction between species selection and species sorting (656–663). All would see a snake-driven contrast in bird lineages as sorting, not selection. But there is an important contrast in views as to how this distinction is to be drawn. On one view (which I have defended), for species selection the differential prospects for speciation and extinction must depend on properties of the species themselves, rather than properties of the individual organisms that compose the species. If one species (or species group) prospered (while another faded) as a result of such factors as (i) the geographic range of the species, (ii) its diverse gene pool, (iii) its division into many subpopulations, (iv) its presence in ecological communities of many different kinds, then (granted a drift/selection distinction) that would indeed be an example of species selection. Such selection could be iterated and hence responsible for shaping lineages, if (for example) species with broad ecological tolerances tended to give rise to similarly tolerant species, and conversely. There is no reason to doubt that species-level properties are significant in determining the fate of species. But often they may not be transmitted to daughter species. Moreover, such properties are difficult to identify in the fossil record. For a single bedding plane rarely represents a moment in time and space, a slice through a functioning community or landscape. It is the aggregated detritus of many communities. Even when species selection has been important—and there are a few very plausible examples—its paleobiological signal may be faint. So on this view species selection is probably real. But it is hard to demonstrate and perhaps is not very common.

Gould, mostly in work with Elisabeth Lloyd, has defended a more inclusive criterion (663–670). We have species selection when the fitness of the species is not a simple function of the fitness of the individuals within the species. The intuitive picture is that we have species selection when the fate of the species is not explained by the extent to which its individual members are adapted. But this criterion turns out to be too inclusive; more inclusive than Gould himself wants. In almost every case, the speciation/extinction propensity of a species will stand in a very complex relationship to the fitness of the individual organisms within the species. This follows from the fact that species are typically divided into many subpopulations, in different communities and conditions, subject to differing selection regimes. The mosaic character of species implies that species fitness is standardly "emergent." It is almost never a simple function of the fitness of individual organisms. Hence, extinction and speciation would almost always, on this criterion, be species selection not species sorting. In short, while it is likely that species selection is a real phenomenon, problems remain with both its empirical vindication and theoretical clarification.

3. The Eye and the Embryo. Gould argues that evolutionary biology has underestimated the importance of developmental biology. His central idea (shared with many others) is that history and development affect evolutionary lineages by biasing the supply of variation to selection. To see his point, it helps to conceptualize the evolution of a lineage as a trajectory through phenotype space; a space whose dimensions capture possible variations in morphology, physiology and behavior. This trajectory might well be affected by the supply of variation to selection. To see why, consider an idealized conception of the production of variation. The power of selection to change the phenotype of a population in any direction is at its maximum if variation around the current phenotype is (i) symmetric: in each dimension, variation occurs with equal frequency on each side of the current mean; (ii) multidimensional: variation occurs in many dimensions of phenotype space; (iii) continuous: in an given dimension, there are many variants with small differences from one another and from the current phenotype; (iv) independent: variation in one phenotype dimension is not tied to variation in an other; (v) iterating: if the average phenotype of the population shifts, the mechanisms of variation produce symmetric, continuous, multidimensional, independent variation around that new mean (1027–1028).

Variation is a precondition of selection. But to the extent that the mechanisms of variation satisfy these conditions, we can treat them as providing only a fuel for evolution. Without these mechanisms, there would be no evolutionary change. But they impart no direction to evolutionary trajectories. The specific path of a lineage through phenotype space will depend on selection (and perhaps drift) but not the mechanisms of variation. To the extent that variation is not symmetrical, continuous, multidimensional, independent, and iterating, the supply of variation will play a role in determining evolutionary trajectories. No one thinks that the supply of variation is perfectly unbiased. Development has evolved, and its evolution makes some changes more likely than others. Differences in view concern the extent and manner to which biases in the supply of variation influence evolutionary trajectories. On Gould's view, these effects are profound and far-reaching. He discusses two very different kinds of effect: *walls* (1028–1032) and *bridges* (somewhat akin to Dennett's cranes) in phenotype space (1033–1037).

In most discussion of historical and developmental constraints on evolution, the focus has been on walls. Certain regions of phenotype space are walled off. Six-legged chordates are unevolvable because the requisite variation is not generated. Gould thinks that such developmental walls have played an important role in the history of animal life. I am not persuaded. Gould thinks the conservatism of animal evolution—the stability of major body plans, and the relatively limited amounts of phenotype space explored in the 550-plus million years of Metazoan evolution—need special explanation. Thus Gould writes of the Metazoan's "skewed and partial occupancy of . . . morphospace" (1056). But how much phenotype space should we expect the Metazoa to occupy, supposing that variation were symmetrical, continuous, multidimensional, independent, and iterating? We have no rigorous way of answering that question, hence no clear way of showing that there is a phenomenon—restricted disparity—to explain.

Suppose this challenge were met, and we knew we ought to be surprised (say) by the lack of six-legged elephants. Perhaps these restrictions in disparity are the result of *fitness trenches.* For sure: six-legged elephants, with their legs fully functional and properly integrated into elephant skeletal and muscular systems, would be biomechanical improvements over our paltry, anorexic four-legged kind. But would intermediate models with rather rudimentary control and attachment arrangements be an improvement? There can be selection-based constraints on the exploration of phenotype space. For the representatives of every selection-driven move in that space must be fitter than their immediate ancestors.

In his case for the importance of developmental walls, Gould lays considerable stress on the most extraordinary discovery of evolutionary molecular biology: the existence of numerous, ancient, widely distributed genetic homologies (1089–1120). For example, to everyone's astonishment, it turned out that genetic homologies underlie the development of fruitfly, cephalopod, and vertebrate eyes. This is one example of many. The HOX gene complex, controlling segment differentiation in fruitflies, has homologies across many metazoan phyla, including the chordates. In vertebrates, these genes are involved in hindbrain development, showing a similar geometric pattern of front-to-back activation to that first found in fruitflies.

These results are wonderful and astonishing. But what do they mean for the flow of variation? They do hint at unexpected conservatism: genetic resources, sometimes modified, sometimes multiplied, are conserved over long periods. But these preserved genetic resources are not tightly coupled to specific phenotypic structures. After all, Pax-6 and its homologies are involved in structures as different as fruitfly and human eyes. The HOX genes involved in the development of our hindbrains tell fruitfly segments whether to sprout wings, halteres, or antenna. So we can read these results very differently, as showing the flexibility of developmental mechanisms. Conserved genetic resources are used to build an extraordinary range of phenotypes. I am sure there are developmental walls, but these deep homologies are not evidence that they have had pervasive effects on the history of life.

The case for the evolutionary importance of developmental bridges is considerably stronger. The idea here is that developmental innovations make certain regions of phenotype space particularly accessible for every species that inherits that invention (1123–1132). Not every animal with a Pax-6 homolog has eyes. But eyes evolve easily in all lineages inheriting this developmental resource. This idea is becoming orthodox in evolutionary developmental biology. Developmental modules—for example, the HOX system of controlling differentiation along a front-to-back axis—are preserved, inherited, sometimes duplicated, and reused. These flexible, multiple-use developmental resources open faster routes into phenotype space. They increase the evolutionary potential of lineages with them, but not equally in all directions. Eyes are common items of perceptual equipment and have evolved in many lineages. Vision is much more common than (say) electrolocation, largely confined to fish. This might be because vision is useful in a larger range of lifestyles, or because the evolution of eyes is not blocked by a fitness trench. There is a smooth array of increasingly efficient eyes, each minimally superior to its predecessor. But it may well be because one key genetic resource for making eyes was invented early and distributed widely. This example is conjecture. But my hunch (with Gould and many others) is that the evolutionary timing and hence the phylogenetic distribution of key developmental innovations is of profound importance to life's history.

4. All Warts, or Warts and All? In many ways this lumbering book is enormously irritating. It is grossly unfair to its intellectual opposition. Inevitably long, it is encrusted with asides, digressions, polemics, feuds, autobiographical excursions, and hobbies. It is in desperate need of the editorial chainsaw. Yet its impressive features (in the end) much outweigh its irritations. For Gould articulates and defends a distinctive vision of the agenda of evolutionary biology, of the mechanisms of evolutionary change, and of the relationship of evolutionary biology to its own past. It is a book of great power, scope and learning. In every sense, it is a fitting monument to its author.