

## Past Present, and Prospective Problems in the Philosophy of Biology

### Abstract

Problems in the philosophy of biology go back to Spinoza, if not to Aristotle. They proliferate after Darwin and become pressing in the late 20<sup>th</sup> century revolution in biology. And these problems are as much those of biologists and philosophers, intersecting as they do both subjects equally. We trace most of these problems back to the nature of adaptation and fitness, and forward in their ramifications for molecular and evolutionary biology and the relations among them. We apply solutions to these problems to illuminate the structure and explanatory strategy of the theory of natural selection, and the controversial conceptual issues about units and levels of biological organization in which it has figured. The philosophy of biology has been a subject of excitement and ferment for more than a generation. We anticipate no relief from this state of affairs.

1. Vitalism and functionalism
2. Fitness, falsifiability and the nature of biological laws
3. Macromolecular reductionism
4. genic reductionism and the levels of selection
5. Conclusion: The agenda of the philosophy of biology

Like biology, the philosophy of biology has been in almost perpetual ferment in the last generation. And the reason is that biology itself has experienced revolutionary development, both in its molecular and evolutionary components. Indeed, the developments in these two domains of biology, evolutionary and molecular have fed one another. Our understanding both of the history of evolution on this planet, and the macromolecular phenomena that ubiquitously underlie present biological activity have heavily informed one another. What these developments suggest is that the mechanism of blind variation and natural selection which Darwin hit upon more than 150 years ago is the most important and perhaps the sole mechanism we require to account for all phenomena above the level of purely chemical processes. What distinguishes molecular biology from organic chemistry, what accounts for the biological organization which is revealed in studies of developmental ontogeny, indeed, the mechanism of learning, the organization of behavior and the character of social institutions, may all turn out to be explainable by this same process of blind variation and selective retention

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

(to employ D.T. Campbell's [1964] term for Darwinian evolution). It is nowadays impossible to confidently identify limits to the ultimate explanatory domain of the theory of natural selection.

As Dobzhansky [1973] once said, and philosophers of biology now intone almost as a mantra, "Nothing in biology makes sense except in the light of evolution." And yet the one thing on which neither biologists nor philosophers of biology have yet secured consensus is the character of evolutionary theory, the meaning of its explanatory concepts, the structure of its theory, and the logic of its explanations. In this chapter we begin by illustrating how ubiquitous the theory is for biology, and then provide an account of the conceptual problems it raises along with our own approach to resolving them. We conclude with an intimation of where the revolution in biology can be expected to carry its philosophy over the medium term future.

### 1. Vitalism and functionalism

Perhaps the most obvious difference between the biological realm and the physical realm is the apparent purposiveness of the former. Such a distinction would not have been made by Aristotle for whom teleological explanations were appropriate in both the physical and the biological realms. (It is worth noting that Lennox [1992] has argued that Aristotle does make a fundamental distinction between the biological and physical, reserving "full-blown" teleology for the biological.)

Aristotle distinguished final causes, goals, purposes, ends, later stages which explained the processes that bring them about, from efficient causes, earlier states which explain the things they bring about. But since Newton such teleological explanations have been otiose in physical science. Elsewhere they are problematical, for as Spinoza argued, explanation in terms of future purposes, as yet unattained goals, and ends that might never be realized, is "to reverse the order of nature" and make the effect the cause. Post-Newtonian mechanists came to regard purposive explanation as a dead metaphor and the vestigial remains of an anthropomorphic view of nature.

Yet such teleological explanation is indispensable in biology. Why does the heart beat? Answer: In order to circulate the blood. This is the correct answer to the question, it actually explains why the heart beats, and it is one of the most important discoveries of the 17th century. Any account of science must accommodate the apparent teleology of these facts. Indeed, for some matters may be more serious if, as many suppose, the very terms in which the biological is described can only be understood teleologically. To call something a wing, a feather, a tissue, a cell, an organelle, a gene, is at least implicitly to describe it in terms of its function, i.e. the purpose it serves in the behavioral

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

economy of some larger system.

It must be kept in mind that biological structures can also be identified independently of their current function. Indeed the possibility of such identification is required for one of the central concepts of evolutionary biology—that of homology. Darwin (1859, 1871) clearly recognized that the most compelling sort of evidence he could marshal for his theory of descent with modifications was the existence of homologues, structures that are similar because of common descent. But homologues are not always easy to distinguish from analogues, structures with independent ancestors that are similar due to common selective pressures. Thus homologues that have either lost their function entirely (rudiments or vestigial organs), or have switched their function were prime evidence for Darwin's theory. "The homological construction of the whole frame in the members of the same class is intelligible, if we admit their descent from a common progenitor, together with their subsequent adaptation to diversified conditions. On any other view the similarity of pattern between the hand of a man or monkey, the foot of a horse, the flipper of a seal, the wing of a bat, &c, is utterly inexplicable."

(Darwin 1871, p. 31). Today, molecular evolutionists are in as much need of the concept of homology as was Darwin, though now it is applied at the level of DNA sequences. Just as we can recognize the wing of a bat and the hand of a human as homologues, we can now recognize developmental regulatory genes in humans (e.g., **PAX6**) as homologous to genes in fruit flies (**eyeless**), though their functions are quite different. Even independent of considerations of homology, anatomists have long identified structural units without knowledge of current function. The lateral line in fish was recognized as an anatomically complex and taxonomically widespread feature long before its function in audition was understood. We should distinguish the functional notion of wing (a device for flight) from the phylogenetic notion of avian wing. Penguins have the later, but not the former. Butterflies have the former, but not the later. Thus it would be mistaken to claim that all biological concepts are functional or teleological. Notwithstanding, biological language is chock-full of teleology and function: even where evolutionary homologues for example do not reflect purposes, they reflect former purposes, and where systems are individuated structurally, the search for their purposes is always on the agenda.

The conclusion to which these considerations led prominent biologists in the 19th century was that, since purpose could not be understood in terms of efficient prior causes, the biological is distinct and autonomous from the physical, and not reducible to it. Despite the mechanists' conviction that the biological was "nothing but the physical" (though more complex by orders of magnitude than

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

an any machine we could construct and understand causally), mechanists had no well supported account of exactly how the purposive even could be causal. The 19<sup>th</sup> century opponents of mechanism had no such trouble. Some invoked the desires and designs of the deity of theism. Those who declined to appeal to god in matters scientific sought the source of the teleological character of biological systems in 'life' or 'vitality', a non-physical property or substance inaccessible to any direct experimental study (Dreisch [1914]).

By the turn of the 20<sup>th</sup> century "vitalism" was thoroughly repudiated by mainstream biology. Nor did it sit well with the empiricism of positivist and post-positivist philosophy of science. In particular it did not fit positivism's vision of science as a hierarchy from biology through chemistry to physics in which reduction characterized the relations among theories and things. Accordingly, 20<sup>th</sup> century philosophers of science sought an analysis of teleological explanation and functional description which freed it from any commitment to the operation of final or any sort of non-physical causes.

The earliest inspiration for causal analyses of purposive behavior was provided by the feedback and feedforward causal loops of cybernetic theory. Taking inspiration from servomechanical devices like the Watt steam engine governor, early post-world war two rocket guidance systems, and bimetallic thermostats, philosophers analyzed teleology as the maintenance of a system's equilibrium state by causal feedback (and feedforward) loops between subsystems as their states responded to whole system's deviations from equilibrium (Nagel, [1961], chapter twelve, Sommerhoff, W. [1950]). The philosophical objective, as with most programs in positivist and post-positivist philosophy of science, was to provide a set of individually necessary and jointly sufficient casual conditions for the truth of "The function of **A** in system **S** is to enable **S** to attain goal **G**". And the criterion of adequacy for such an analysis was the absence of two kinds of counterexamples: those in which all conditions are filled, but no function seems to be realized, and those in which some conditions are not satisfied and yet a function is being fulfilled. Thus, one sort of objection to the definition of function in terms of feedback/feedforward loops, would be the claim obvious to "common sense" that it is not the function of green plants to maintain an atmosphere with a certain ratio of oxygen to CO<sub>2</sub>, even though they are part of a system composed of plant material, CO<sub>2</sub>, soil, water, which (until recently) did so. Another purported counterexample involved functions which fail to secure their goals: again, our "linguistic intuitions" tell us it will still be the function of chlorophyll to catalyze the production of starch from CO<sub>2</sub> and water, even when owing to the absence of water it does not do so.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

By the mid-1970s the succession of analyses of functional attribution in terms of feed back/feed forward systems eventuated in Robert Cummins' [1975] treatment of functions as "nested capacities", a treatment which has come to be termed the causal role (CR) analysis. On this view,  $x$  functions as an  $F$  in some system  $s$ , "relative to some analytical account" of  $s$ 's capacity to  $G$ , just in case  $x$  is capable of  $F$ ing in  $s$  and the analytical account appropriately explains  $s$ 's capacity to  $G$  in part by appeal to the capacity of  $x$  to  $F$  in  $s$ . There are thus two sets of capacities involved in a functional attribution, the wider capacities of the whole system, and the narrower one of the component function which enable the whole system to discharge its wider capacities. Cummins writes that such an analysis "allows us to see exercises of the analyzed capacity as programmed exercise of the analyzing capacities." The CR analysis appeals to complex concepts such as "programmed exercise", and "analytical account", as well as crucial distinctions such as wider/narrower or more complex/simpler capacities. The appeal to these notions on the one hand enables Cummins' account to avoid counterexamples of the sort originally advanced against feedback/feed forward accounts, but on the other hand this appeal may threaten the account with vacuity. For example, if "programming" is itself a teleological affair, the nested capacities analysis is question-begging. Moreover, if the very taxonomy employed to individuate systems of nested capacities requires a prior identification of function or purpose, the CR account will be circular or presuppose a more fundamental analysis.

One cogent defense against these charges of vacuity or regress focuses on the role of CR functions, as nested capacities—in biological disciplines such as comparative anatomy. Amundson and Lauder [1998] give a version of the CR model as deployed to account for the operation of a type of complex system which can be identified morphologically (e.g., the "forelimb"), even when instances of such a type have diverse functions (e.g. leg, wing, fin, flipper). A CR account of the capacities of such a system cashes them in for the capacities of some of its parts also identified morphologically (e.g., the ulna), and not in terms of their purposes. This sort of functional explanation may be said to provide an account of how things work, not why they work the way that they do. Such explanation is a necessary part of biology, but as even its defenders recognize, this sense of function is not the one at issue between mechanists and their opponents (See Amundsen and Lauder, 1998], p. 354). It is "function without purpose".

The cybernetic and the CR approach never fully satisfied philosophers struggling to assimilate purpose to causation. And in the long period during which philosophers have sought a causal account of teleology, the fashions and standards of philosophical analysis changed. In particular, philos

philosophers came eventually to pay much more attention to the actual role of concept like "function" in biology, and less attention to what unschooled "linguistic intuitions" suggest about these concepts. Philosophers came early to appreciate that the purpose or function of a structure, an organ, a process, or a behavior had something to do with adaptation. Since the biological account of adaptation is Darwinian, it was concluded that the theory of natural selection must underlie any account of biological purpose or teleology.

Darwin must have early seen the relevance of natural selection to function. He wrote: "We hardly use the word ought in a metaphorical sense when we say hounds ought to hunt". His point must have been that it is literally true that the function of the hound is to hunt just in case it was selected for hunting.. A century later Ernst Mayr [1961] drew the now widely accepted distinction between proximate and ultimate explanations in biology. Proximate explanations, explain how some extant biological system works, while ultimate explanations explain why they emerged over time as the result of a history of evolution by natural selection. Thus purposive functions must be the result of the latter process while CR functions are invoked in proximate explanations.

Larry Wright [1976] was the first philosopher to develop this sort of insight into a general analysis of teleology—the so-called etiological or selected effect (SE) analysis. According to Wright, 'the function of  $x$  is  $F$ ' means a)  $x$  is there because it does  $F$ , b)  $F$  is a consequence of  $x$ 's being there. The apparent circularity is dissipated once we introduce a type/token distinction. 'Heart' names the general category or type of which our hearts, and those of all other mammals, are instances or tokens. We are to understand Wright's definition as a claim about tokens and types: tokens of type  $x$  have the function of  $F$  only if other tokens of the type  $x$  have existed in the past because they have had  $F$  as one of their effect.

Wright was concerned with a general analysis of teleology, one not tied down to any specific biological process or theory. In contrast, and quite properly in our view, most philosophers of biology have sought to apply this analysis of teleology to biology, and so have tied it to a specific biological mechanism—evolution by natural selection. Thus, the (SE) function of  $x$  in an organism  $O$  is that effect of past instances of  $x$  that increased the fitness of the bearers of  $x$  relative to other organisms within the lineage bearing variants of  $x$ . Equivalently, the function of  $x$  in  $O$  is  $F$  if and only if  $x$  is an adaptation (in the lineage of which  $O$  is a member) for  $F$ . For example, if the SE function of wings in birds is flight, then wings are adaptations in birds for flight. The purely causal—mechanical explanation of the evolution of an adaptation in a

lineage is the relevant etiology in explaining why the trait exists now (in the form and frequency it does exhibit). Abstractly,  $x$ 's existence is explained in terms of the effects of past instances of  $x$ ; but not just any effects: we cite only those effects relevant to the adaptedness of possessors of  $x$ . Whence the label 'selected effect' analysis, which both reminds us of the role of natural selection as the relevant etiology, and highlights the fact that functions are a special kind of effect.

Philosophy has come as close to consensus as it is ever likely to get on the adequacy of the SE account of functional attribution and purposive explanation. But two questions remain, one narrow, and one broader: (1) What is the relationship between the CR account and the SE account? And (2) Does the SE account explain away purposes or does it naturalize them? As for (1), our view is that the two analyses are complementary. Any complete, or nearly complete adaptational explanation of the existence of a trait in terms of its function, will contain a CR analysis (Brandon [1990], and Rosenberg [2000]). If bird wings are indeed adaptations for flight, then any complete explanation of the evolution of wings will include a causal role analysis of the aerodynamic effects of the shape and motion of bird wings, showing how their capacities lead to the adaptation they realize, the capacity to fly. Conversely, many of the most interesting CR explanations in biology will be explanations of traits independently identified as adaptations. In this way Cummins' original reference to programming is cashed in for non-teleological, purely causal factors of random variation and natural selection.

Question (2) poses the contrast between explaining purpose as a complex sort of efficient causality as opposed to explaining it **away** as an illusory overlay which the human mind casts over complex but purely causal processes. Darwin's followers have not agreed about whether the theory of natural selection explained purposes or explained them away. Following Asa Gray, Thomas Huxley, Darwin's self appointed "bulldog", wrote that "the most remarkable service to the philosophy of Biology rendered by Mr. Darwin is the reconciliation of Teleology and Morphology, and the explanation of the facts of both, which his view offers" On the other hand, Michael Ghiselin ([1969], p. 136) insists

...a myth has grown up, partly through the work of [Asa] Gray, partly through the work of Darwin's son and biographer, Frances Darwin, that Darwin somehow "brought teleology back into biology." In any nontrivial sense of that word, he did the exact opposite, getting rid of teleology and replacing it with a new way of thinking about adaptation.

In most respects there is nothing more than a verbal difference between explaining away tel

eology in biology and making purpose safe for a mechanistic view of nature. What is not a merely verbal dispute is the issue of whether Darwin's achievement constituted the completion of the scientific revolution that began with Galileo and Kepler. Does the theory of natural selection show that, in biology at least, explanation can dispense with ultimate purposes that somehow transcend the mechanisms of evolution by natural selection? Does the theory show that there is no scientific basis for the conclusion that nature writ large has a transcendent purpose or meaning that makes it ultimately "intelligible"? We think that it does. Does the process of natural selection allow for the interpretation of evolution as globally progressive, or for the notion that human kind represents a higher, more advanced outcome of natural selection than any other species? We know that it does not.

## 2. Fitness, falsifiability and the nature of biological laws.

The selected effects analysis of teleology cashes in biology's claims about purpose, goal, and end for the sort of causal relations physics and chemistry have taught us to know and understand. However, the analysis leaves one very large hostage to philosophical fortune. Its reliance on the theory of natural selection shifted the focus of the mechanist's traditional suspicion from teleology to Darwinian theory. Those skeptical about the theory of natural selection on the grounds that it is very different from theories in physical science made common cause with those whose religious motives led them to cast doubt on Darwinism's credentials out of a desire to protect the biological and the human sphere from absorption by physical science. Both parties long argued that the theory of natural selection was an empty tautology, that its claims were unfalsifiable pseudo-science, and that its explanatory concepts could not be given the sort of empirical grounding that real science requires. Some still so argue. These charges are potentially serious: they would for example, unravel the consensus on teleology as selected effects. Accordingly, much work in the first two decades of the post war development of the philosophy of biology focused on dealing with them.

The problem is a relatively easy one to state. In the slogan Darwinians adopted from Herbert Spencer, evolution is the result of the survival of the fittest. Or rather, it is the consequence of differences in fitness resulting in greater reproduction, and along with persistent variation in the heritable traits that make for fitness differences. These produce descent with modification: i.e. evolution. But, what exactly are fitness differences? If we answer, as many biologists have incautiously done (e.g. Simpson, [1949]), that fitness differences are defined in terms of differences in reproduction, the slogan becomes the triviality that evolution is the differential reproduction of those organisms which

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

h reproduce differentially. Such a claim is a necessary truth but an empty one, the direct consequence of definitions, an "analytic sentence" or tautology, which can have no explanatory power (compare, does "All bachelors are unmarried adult males Homo sapiens" explain why a particular bachelor, say, Elton John, is an unmarried adult male Homo sapiens?). Nor can such a claim be falsified by any logically possible observations. It is for this reason that Karl Popper initially stigmatized Darwinian theory as unscientific. [Popper, "Autobiography", *The Philosophy of Karl Popper*, ed Schillp, 1974, p. 1340). Recognizing the importance of the theory to biology, Popper came eventually to describe Darwinism as metaphysical research program. Even under this more generous construal, the prospects for the theory of natural selection's passing muster as a scientific theory are reduced. And this is often enough of a basis on which religiously motivated opponents of Darwinism have argued that, like that other metaphysical doctrine, theism, it should not be taught in schools as science.

The "tautology problem", as it has come to be called, arises when one tries to state in a declarative form the core thesis of Darwin's theory, some Principle of Natural Selection (hereafter, PNS). One could simply deny that there is a core thesis in Darwin's theory, deny that there is a PNS. For example, Beatty ([1992], p. 118) observes that "If some "principle of natural selection" were indeed the evolutionary equivalent of Newton's laws of motion then one would expect to find it prominently placed in the standard textbooks; but it is not." Though little weight should be attached to the evidence of textbooks alone, it remains true that denying that there is a PNS will enable one to escape the tautology problem. But it is then incumbent on someone pursuing this strategy to state just what is central to the theory. Some philosophers have taken up this challenge. Focusing on the fact that biologists frequently describe evolution as change in gene frequencies, some philosophers (e.g. Ruse, [1973], Lloyd [1989], Thompson [1988], Sober [1993]) have claimed that models from population genetics form the core of neo-Darwinian theory. Their view is often closely associated with a so-called semantic approach to theories.

The traditional view of theories holds them to be sets of (approximations to) general laws which work together, often in a way that can be axiomatized, to explain non-strict empirical generalizations, and their exceptions, and to predict new general regularities, and do all three with accuracy that increases as the degree of approximation to strict laws improves. Embracing this traditional view makes it imperative that we identify a PNS as one of the laws which constitute this theory. But, the philosophers of biology who deny the need for a PNS in the theory contend that this picture is not borne out in biology, and especially not in evolutionary biology. Here one is confronted by mathe

mathematical models, such as the Hardy-Weinberg equation, or the Mendelian model of segregation and independent assortment, or the Fisher sex-ratio model, etc. These models are in effect implicit definitions of the systems of heredity they describe. The question which evolutionary biologists confront is which among the set of models at the biologist's disposition are appropriate for systematizing the varying sets of data which the biologist confronts. The answer to this question will of course vary given the objectives and the data of the scientist. On the semantic view, the theory of evolution is this set of models.

But this approach is problematic in two ways. First, merely asserting that the theory of natural selection is constituted by a set of models which share some overlapping features, does not really obviate the need for a PNS. To see this, consider the question, why do just these models and not some others "work" to systematize and explain evolutionary phenomena? The only cogent answer is because they reflect some regularity in nature, which is just what is invoked in a PNS of the sort the semantic approach seeks to avoid. Second, the fact that biological processes satisfy these models well enough for them to have explanatory and predictive relevance to nature is the consequence of, not the explanation of natural selection. As Beatty showed [1981], models in population biology such as the so-called "Hardy-Weinberg law" about the constancy of gene frequencies in the absence of selective forces, could not possibly describe processes fundamental to a general evolutionary theory, for that theory must apply to evolution during the first billion or so years of life on Earth, long before genes existed. Diploidy, and sex, which are reflected by the models in question, are highly evolved mechanisms. Accordingly, a general theory of evolution should be expected to explain their evolution, as opposed to presupposing their prior existence.

It may be controversial whether, as we hold, biology should, like physics, aim for theories that are true everywhere and always in the universe. But surely it is patent that a theory of evolutionary biology should account for the evolution of life on Earth. And indeed, virtually all contemporary evolutionary biologists would accept this demand. Accordingly, they require a PNS, whether they explicitly advance one or not. Let us see if we can construct one that does justice to Darwinism as a general theory of natural selection true everywhere and always.

Consider the following statement of the PNS: If **a** is fitter than **b** in environment **E**, then (probably) **a** will out-reproduce **b** in **E**. The question then becomes: How can we define relative fitness (or relative adaptedness) in a way that makes PNS central to Darwinian theory and the explanations it provides? Logically there are only three possible approaches. The first is to simply define the

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

fitness of a biological entity in terms of its actual reproductive success. That not only makes the P/NS a tautology, it makes it an overly cautious statement of a tautology. The 'probably' could safely be deleted. (Only the most cautious would say "Bachelors are probably unmarried" .)

In contrast, the second option seems most attractive from the point of view of most standard theories of explanation. That is, to define 'fitness' in terms of some definite set of biological (or chemical, or physical) properties, such that organisms with that set of properties are inevitably selected over competitors with different sets of properties. For instance, it has been suggested that energetic efficiency is such a property (Bock and von Walther [1965]). One of the fundamental challenges of life is to turn environmental energy into usable energy. Accordingly, it might be claimed that the more efficiently one does that the better off one is from an evolutionary perspective. Thus, as between two organisms, the fitter is the more energy-efficient one. But that, unfortunately, is often false.

A large aggressive male animal may be able to garner more environmental resources than it could possibly use (and thus fail to be the most energetically efficient) and will be selected for owing to the fact that females treat the conspicuous waste of resources as a mark of the male's likelihood of having many or strong or otherwise fitter offspring. Attempts to circumvent this and other counterexamples will simply narrow the scope of this definition until it becomes useless as a stipulation about fitness in general. But the same fate faces any attempt to define fitness in terms of its causes or prior determinants. The trouble is that different biological (or chemical or physical) properties are selected among different organisms in different selective environments with equivalent results in fitness (cf. Rosenberg [1978] and the discussion of "supervenience" below).

The third option is generally called the propensity interpretation of fitness, and is widely accepted among philosophers and methodologically sensitive biologists. But it must be emphasized at the outset that there are really multiple propensity interpretations of fitness, and substantial disagreements among the proponents of these various approaches. Indeed, there are so many alternatives that come under this label, that it is tempting to delimit this set of approaches by exclusion: any approach to defining fitness (or adaptedness) neither in terms of actual reproductive success nor in terms of some specific biological property or set of such properties will be labeled a propensity interpretation. Some we would so label (e.g., Michod [1985]) would be unhappy with the label. But that is because they disagree with the details of one or more of the other approaches that happily accept the label.

These disagreements are important, and will be briefly explained here, but they should not obscure what is common among propensity approaches, namely the belief that what is required is an abstract

, schematic, probabilistic (as opposed to biological, chemical or physical) definition of fitness. We develop a positive characterization of the general strategy of a probabilistic propensity definition below.

The tautology problem shows us that defining fitness in terms of actual reproductive success will be non-explanatory. Similarly, we cannot explain why the toss of a fair coin yields heads by saying that it was a toss that yielded a head, nor perhaps can we explain why the coin toss yielded heads by adverting to the fact that its probability of doing so was .5. For fairness of a coin is presumably defined as  $\Pr(T) = \Pr(H) = .5$ . But we can provide at least part of the explanation of certain actual outcomes or their features by referring to their probabilities or chances of occurring. Suppose we have two coins, one with a previous record of tosses biased towards tails ( $\Pr(T) = .7$ ) and the other with a record biased towards heads ( $\Pr(H) = .7$ ). Supposed we then toss both coins ten times and record the results: the first coin comes up tails 8 times and the second comes up heads 6 times. Now, it is an empirical regularity, in fact, a law of nature, so evident that we never explicitly articulate it, that, **ceteris paribus** outcomes with higher probabilities in fact obtain more frequently than outcomes with lower probabilities. Following a tradition among students of inductive logic, we may call this the principle of direct inference (or alternatively the law of likelihood), for it allows inference directly from a probability to an outcome. The only possible challenge to this claim would object to its alleged status as an empirical law, and insist that it is in fact a disguised tautology true in virtue of the meaning of the word ‘probability’. However, as philosophers have long recognized, on any available interpretation of probability as either long run relative frequency, or subjective Bayesian betting odds, or pure probabilistic propensity, this proposition is not demonstratively true, and so not an empirically empty tautology. After all, we know that it is perfectly consistent with this principle that a coin with a  $\Pr(T) = .7$  to come up heads one thousand consecutive times. We know this will not happen as a matter of empirical fact, and the reason we do is that we know the empirical, indeed, nomological regularity that the principle of direct inference reports. The principle of direct inference thus gives us good factual reasons to expect that the second coin will yield more heads than the first. Accordingly it identifies factors—probabilistic properties of the coins—distinct from, but explanatorily relevant, to the actual outcomes. The two probabilities in this case of course are presumably determined by asymmetries in the physical composition of the two coins. But they are nevertheless distinct dispositional properties different from those manifest asymmetrical properties on which they rest, and equally they are properties distinct from the outcomes they bring about.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

There remains to the philosopher of science the very real problem of explaining exactly what sort of properties the probabilities in question are. It is a problem with which philosophy of science has self-consciously wrestled at least since John Maynard Keynes' Treatise on Probability [1921]. That there is as yet no solution to the problem of exactly what sort of a property this probability is, does not detract from the reality of probabilities and their empirical explanatory power. It is these same probabilities which the propensity interpretation of fitness invokes. And the difficulties both of identifying the PNS and recognizing its contingent, empirical explanatory power are to be understood once we recognize the invisibility and the empirical character of what we have called the principle of direct inference..

Fitness is a probabilistic propensity. It is a real property of organisms (in their specific environments), in just the way that the chance of heads is a property of a coin (and its tossing device) . That having been said, there still remain the disagreements among philosophers about which among different possible probabilistic propensities to leave descendants best defines and/or measures fitness. Brandon [1978] and Mills and Beatty [1979] were the first to explicitly propose propensity interpretations of fitness . They both defined the propensity as the arithmetic mean of the various possible immediate reproductive outcomes (e.g., having 1, 2, 3, or ... up to  $n$  offspring) weighted by the probability of each of these  $n$  outcomes. Thus they defined fitness as the expected reproductive success (in the standard mathematical meaning of expected value.) But as early as 1973, the biologist John Gillespie [1975] showed that this in some selective circumstances this definition was not generally correct. In certain situations the geometric mean of the various probabilistically weighted  $n$  outcomes was a more appropriate measure of fitness than the arithmetic mean. More generally what Gillespie showed was that the mean of the probability distribution of different reproductive outcomes was one important thing that selection operates on, but that selection also operates on the variance of reproductive outcomes. Increased variances of reproductive outcomes are always selected against. The variance itself can take different forms, it can be temporal, spatial, or stochastic. Selection acts differently depending on the form of the variance.

A simple example illustrates the selective disadvantage of increased variance in reproduction. Suppose a population consists of two types of annual plants **G1** and **G2**. Suppose further that they reproduce asexually and reproduce true to type, i.e., **G1**'s produce **G1**'s and similarly for **G2**'s, and that they have an equal mean reproductive success of 2. But **G1**'s achieve this mean by producing 1 offspring one generation, and 3 offspring the next, while **G2**'s always produce 2 offspring. (

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

Gillespie would label this sort of variance temporal, as opposed to spatial or stochastic. Differences among these types of variance do not affect the basic point.) Start off with  $n$  individuals of both types in generation 1, and suppose this is a good season for **G1**'s, i.e., that they have 3 offspring each. Thus generation 2 has  $3n$  **G1** and only  $2n$  **G2**'s. The  $3n$  **G1**'s now each have 1 offspring, while the  $2n$  **G2**'s each have 2. Thus generation 3 has  $3n$  **G1**'s and  $4n$  **G2**'s. Again it becomes a good year for **G1**'s and so they each produce 3 offspring. Thus generation 4 has  $9n$  **G1**'s and  $8n$  **G2**'s. But this will be the last generation **G1**'s will outnumber **G2**'s. Generation 5 will have  $9n$  **G1**'s and  $16n$  **G2**'s. By generation 10 the **G2**'s will outnumber the **G1**'s by a ratio of 512 to 243. It only gets worse for the **G1**'s. What is happening in this simple example is that the advantage **G1**'s have over **G2**'s in their good seasons, when they produce 3 compared to 2 offspring, is more than offset by the disadvantage of producing 1 compared to 2 in the bad years. The same disadvantage pertains when the variance is distributed over space, rather than generational time, or when it is distributed stochastically.

Bearing Gillespie's arguments in mind Brandon [1990] provides a revised definition of fitness: it is essentially the arithmetic mean of the reproduction rates for some future generation discounted by some (positive) function of the variance of those rates. The exact form of this function is determined by the nature of the variance as it reflects differences in the selecting environment. This definition is thus itself schematic, and must be distinguished from the various units in which fitness can be measured. Once the form of the variance and the later generation number has been specified, the schematic content of the definition is satisfied by these parameters, and the units in which fitness is to be measured on a given occasion are given. It is crucial here as elsewhere in science to distinguish a dimension from the units in which it is measured: time is not identical to milliseconds as measured on some standard clock, space is not identical to centimeters as measured by a standard meter stick. Though each is measured in such units, they are not always accurate and require correction.

Similarly, greater fitness is not identical to any particular probabilistic propensity to have more offspring or a smaller variance in them, or to have more grand-offspring or great grand offspring, or some variance in them either, and so forth. But from case to case it is measured by one or another of these units, depending on the selective environment. Thus when it is argued (for example by Beatty [1992] that Gillespie's results refute the propensity interpretation, what is in fact shown is that no single set of units accurately measures fitness differences in all circumstances. But still some propensity interpretation or other is always appropriate in each setting.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

In the vast majority of cases, of course, expected rates of reproduction in the next generation provide both the easiest and at the same time perfectly adequate units in which to measure relative fitness. This is in part because most environments select for larger numbers of offspring, as opposed to higher “quality” offspring. The appropriateness of a measure which is relatively short-sighted—making immediate numbers of offspring the relevant units—reflects the fact that natural selection has no foresight. It cannot discriminate among competing types today based on, say, the probability of survival for their distant descendants. Thus fitness measures should reflect what selection “sees”. With few exceptions, a single full generation measure of fitness is what is appropriate. But in the end, it is crucial to bear in mind that what is selected for as fitter over a number of generations may be selected against in the near future. Selective environments change and sometimes a new one will select against the very probabilistic propensity to reproduce which the previous environment made the most appropriate measure of fitness.

Important conclusions for the nature of the PNS follow from this understanding of fitness as a schematic probabilistic propensity function. For it turns out that the PNS is in fact a special case of what we have called the principle of direct inference, a case in which the explanatory or determining property is a biological property, the probabilistic propensity to leave a greater number of descendants, and the explained outcome is a biological outcome, the difference in descendants left. As the biological version of the more general principle, the PNS will share with the principle of direct inference the properties of being a contingent empirical **ceteris paribus** law, and also the property of being so obvious and so implicit in its employment as to appear to be vacuous and without explanatory force. We can now understand why the PNS has so often, but nevertheless mistakenly been stigmatized as a definition.

We can say more, on this analysis, about the notion of fitness, qua probabilistic propensity. Its relationship to features of organisms and environments will be the same in kind as the relationship between the probability to land tails on the one hand and the coin and tossing apparatus which bears this probabilistic propensity on the other. But the relationship will be vastly more complex in degree. For presumably the property that determines the probability of landing heads is just the physical asymmetry of the coin and/or the equipment which tosses it, while the list of traits of the organism and features of its environment which together fix a level of fitness is disjunctive and indefinitely long as well. This relationship is one to which we will return in the next section.

In light of the clarification of the PNS as the fundamental law of natural selection, what follows

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

ows about the structure of the theory of natural selection? The absence of other laws in evolutionary biology and elsewhere has potentially serious consequences for biological theory and the explanations it provides, for most accounts identify laws as the source of a theory's explanatory power. As noted above, one response to the alleged absence of laws in evolutionary theory has involved adopting the so called "semantic account" of the nature of theories. The attractions of this view are evident: it accords with the recognized practice of many biologists. They are notoriously reluctant to state any hard and fast evolutionary laws. Even in classical genetics, when faced with phenomena such as crossover, linkage, meiotic drive, and other exceptions to Mendelian inheritance, biologists make no attempt to refine, improve, or correct Mendel's original model. Indeed, it continues to be employed where appropriate, i.e. when its contrary-to-fact assumptions are harmless for predictive or explanatory purposes (Thompson, [1988], Lloyd, [1989]). Citing Darwin, some biologists and philosophers insist that the theory of natural selection must be understood as a historical claim about the diversity and complexity of biological phenomena on the Earth, and not as a set of generalizations which might for all we know have instances elsewhere in the universe (Sober, [1993]). On this view, explanation in biology is historical not nomological, and does not require laws; models will suffice. Freed from any commitment to such laws, the semantic view of theories is much more easily accommodated to this conception.

However, the absence of underlying laws to provide for the explanatory power of such models would require an entirely new conception of explanation in biology. Our understanding of the PNS makes such an account of biological theory and explanation unnecessary. The models which a semantic approach treats as the theory are better understood as schematic mathematical models, whose empirical specification results in empirical generalizations that are ultimately themselves to be explained by the PNS. As noted above, it will be no surprise that the PNS is almost invisible in biological theory and the explanations it provides. The principle is an obvious consequence of a proposition all accept and many suppose to be a truism, the principle of direct inference. Consequently it is a surprising source of the nomological "glue" that holds disparate evolutionary outcomes together as the contingent consequences of the operation of the PNS on local boundary conditions over the 4 billion years of the Earth's existence, and holds equally well wherever and whenever in the universe environments differentially favor reproducing variants, no matter what their composition.

The philosophical problem that this understanding of fitness and the PNS leaves behind is an enduring question which philosophers have faced at least since the 19<sup>th</sup> century revolution in their

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

modynamics. This is the question, which evolutionary biology shares with thermodynamics, of exactly how to understand the probabilities the two theories trade in. They cannot be the brute unexplained probabilistic propensities of quantum mechanics, for they are both to be explained by the underlying non-quantum processes which give rise to them. Nor can they be assimilated to the [Bayesian] subjective degree of belief that characterizes rational uncertainty in decision-making. To say, the probability of PNS or the second law of thermodynamics is sui generis and different from either is to identify a problem, not to solve it. But the problem is one for philosophers, and perhaps physicists, but not biologists.

### 3. Macromolecular reductionism

Surely molecular biology can escape the problems that daunt evolutionary biology. Here at last one might have supposed that matters have been settled by the advance of biochemistry, and in particular by the macromolecular elucidation of biological processes like heredity, respiration, immunity, and most recently embryological development. This hope, to escape philosophical heavy weather, turns out to be in vain. But the understanding of intertheoretical relations the philosophy of science can provide here also advances our understanding of biology, its epistemology and metaphysics.

It is the reduction of classical to molecular genetics that seems furthest advanced and offers the best case for reduction in biology. Every one agrees that Watson and Crick revealed that the gene is nothing but a nucleic acid molecule, and that the central dogma of DNA transcription to RNA, and RNA translation to protein provides the non-biological account of the chemical constituents which implement the genotype/phenotype relation that Mendelian genetics rests upon.

Genes are nothing but DNA molecules, and yet when we attempt to apply the apparatus of theory-reduction drawn from the physical sciences to biology, no part of it seems to fit. This account of reduction as a form of intertheoretical explanation has two entirely independent requirements: If the narrower, less general, less fundamental, reduced theory is  $T'$  and the broader, more general, more fundamental, reducing theory is  $T$ , then a) The basic underived laws of  $T'$  must be derived from the laws of  $T$  and b) the distinctive theoretical vocabulary of  $T'$  must be linked to that of  $T$ , so that the derivation of the laws of  $T'$ , expressed in its distinctive vocabulary, from those of  $T$ , expressed in its distinctive vocabulary assures that  $T$  explains  $T'$ . To recur to an example in section 2 above, the reduction of the ideal gas law  $PV = nRT$  to the kinetic theory of gases required that 'temperature', a property in the gas law, be set equal to 'mean kinetic energy of the constituent mole

cules', a property described by the reducing theory.

Satisfying requirement a) is difficult enough, for doing so involves establishing that Mendel's "laws" are indeed nomological generalizations or sufficiently close approximations to them for present purposes, and that there are laws in molecular biology about sufficiently long chains of deoxyribonucleases from which these Mendelian principles may be derived. This requirement of laws will be hard to satisfy, as we have already seen. But more attention has been paid to attempts to satisfy requirement b) for on the one hand, doing so is a necessary condition for satisfying a), and on the other hand, it was anticipated that it would be relatively easy to show that the 'gene' of classical genetics is identical to the 'gene' of molecular biology. Alas, these expectations were defeated. To begin with, the classical gene is a unit of linkage, recombination, mutation and phenotypic control. But molecular genetics distributes these functions to different amounts of DNA often physically separated from one another. The single nucleotide is the unit of mutation; breakage and annealing of a DNA molecule can produce linkage and recombination within genes as well as between them; and even on the arguable assumption that a phenotype is a single protein product, no molecular gene sequence is itself causally sufficient to produce the protein, many other genes and their products are also required. The classical gene/molecular gene relationship is not a one-to-one, but a many-to-many relation which precludes the sort of derivation reductionism requires.

Non-molecular biologists and even some molecular biologists might welcome this conclusion because it vindicates their view that many non-molecular explanations in biology are perfectly adequate, complete and indeed unlikely to be improved, still less superseded by macromolecular explanations. But this is a troubling conclusion, for it seems to be at variance with biological consensus that classical genes just are nothing but molecular genes, and the agreement that molecular genetics can (at least in part) shed light on aspects of the processes in which classical genes figure. Philosophers of biology have strongly felt the urgency to reconcile these apparently discordant threads in biologist's explanatory practice.

First, in order to explain the sense in which classical genes are nothing but nucleic acids, philosophers invoked the "type/token" distinction. 'Types' are kinds, while tokens are their instances. Thus, 'clock' is a kind, while 'Big Ben,' 'The Sun' and 'the watch I am now wearing on my right wrist' name three tokens of the kind 'clock.' Notice that 'clock' is a functional kind, that is one whose instances may all be identified by a single common causal roles, in this case, the causal role of time keeping. They do not all play this role equally well, nor do they do so by implementation of the

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

same mechanism. The Sun can be used reliably to measure days, Big Ben can be employed accurately to measure days, hours and minutes, while my watch measures to tenths of a second. And of course these three tokens, like other tokens of the type ‘clock’—tree rings, leaves on trees, the moon, the tides, pendula, hourglasses, atomic clocks, etc.—function through the operation of many different physical systems of parts composed of many different substances. This makes the description ‘ $x$  is a clock’ one which is “substrate-neutral”. Clocks can be composed of an indefinite number of different things arranged in an indefinite number of different ways, just so long as this diversity of arrangements all fill the same causal role—i.e. enable us to tell time—to count equal intervals of it. This means that (unlike ‘temperature’) ‘ $x$  is a clock’ is a type that cannot be defined in terms of or reduced to some finitely enumerable list of structural or physical types. The type ‘ $x$  is a clock’ is not identical to any physical type. Yet there is nothing more to being a clock than to be a purely physical arrangement. Surely every token of a clock is nothing but the token of a physical system. The way philosophers have hit upon to express token-identity without type-identity appeal to the notion of supervenience (Rosenberg, [1978], [1985]): Every token of type **T**’ is numerically identical to some token of type **T**’\* just in case whenever token **a** of **T**’ is nothing but a token of **T**’\*, anything else that is **T**’-type identical to **a**, is **T**’\*-type identical to **a**. Being a clock supervenes on physical structure: even though we can’t give finitely long physical definition of a clock, if anything, no matter what its physical structure, is a clock, then anything else with exactly the same physical structure must also be a clock. Even if there are an infinite number of different ways that some arrangement of matter can be a clock, it will still be the case that being a clock is nothing but being some arrangement of physical matter or other. How does all this help us understand the relationship between the classical gene and the molecular gene?

Distinguish classical gene types and classical gene tokens. A classical gene type is any thing which plays a certain causal role, say, participates in recombination, crossover, linkage, mutation, for some phenotype or phenotypic difference, etc. Now suppose that every particular token of a classical gene-type is nothing but some perhaps quite complicated combination of nucleic acids (including those sequences required for promoters, operons, ribosomal RNA, and all the machinery needed to produce to relevant protein or other phenotype which figures in the classical genetic identification of the classical gene-type in question). As we know, given the redundancy in the genetic code, the existence of multiple reading frames, the neutrality of many base-pair sequences, the existence of repeat sequences, etc. quite different nucleic acid sequences, and sets of sequences can give rise, in th

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

the same molecular milieu, to the same classical phenotype (or phenotypic difference), and so count as tokens of the same classical gene type. But provided classical gene types are supervenient on molecular genotypes, we may hold that classical genes are nothing more than molecular genes even though we cannot reduce the classical gene type to the types of molecular biology.

The supervenience of classical genetics on molecular genetics also holds out the promise of reconciling the claim that classical genes are nothing but macromolecules with the explanatory indispensability of classical genetics. And if successful in this case, the same strategy can be employed to underwrite the autonomy of all the rest of functional biology from molecular biology and from the physical sciences, even while granting the reductionist's claim that the physical facts fix all the biological facts, that the biological is nothing but the physical, and that beyond natural selection there are no forces, entities, processes or properties unfamiliar to physical science to be met with in biology.

The supervenience of classical genes on macromolecules means that for any single process in which classical genes engage, there is an indefinitely large and perhaps motley disjunction of macromolecular processes which "realize", "instantiate", "constitute" this single classical genetical process in question. As such, any reductive macromolecular explanation of why and how this type of process occurs in general will itself have to advert to this indefinitely large disjunction of processes any one of which can and sometimes does bring about the classical genetic processes. But, it is often argued, a) this disjunction of disparate macromolecular realizations of a classical genetic process is not itself explanatory, and b) seeking descriptions and explanations couched in the terms of molecular biology would blind us to important generalizations at higher levels of biological organization (of which the classical genetic level is only one example, Kitcher [1984], [1991]). These two gambits of course require substantial philosophical justification. The argument for a) often proceeds by suggesting that a disjunction of diverse processes cannot be a real, explanatory kind, cannot itself be a (single type of) cause, even though each of its disjuncts brings about the same outcome (Sober, [1984]). Another argument for the same claim presupposes a theory of explanation as "erotetic" i.e. explanations are adequate to the extent they respond to interlocutor's interests and background information. Since some biological interlocutors' interests and backgrounds are not biochemical, a macromolecular reduction, even were it possible, would not be explanatory. Argument b), that we would lose sight of (or leave undiscovered) generalizations that obtain at higher levels were we to substitute the motley crew of macromolecular realizations for the entities and processes that supervene upon

them, of course turns on the claim that there are explanatory laws at these levels (as opposed to the level of the PNS). If these claims can be made out, a resulting reconciliation of physicalism (the physical facts fix all the facts, including the biological ones) and antireductionism (the biological is autonomous from the molecular) will have been effected. This would vindicate a widely shared view among philosophers and biologists.

Regardless of whether the vindication can be established, the notion of supervenience can do much work to illuminate intertheoretical relations in biology. Recall, for instance, the problem of giving a definition of fitness in terms of the factors which determine it, instead of its effects. We compared the probability of a coin's landing heads with an organism's probabilistic expected reproductive rate. The former is simply and directly, indeed reductively explained in terms of the asymmetry of the coin. The latter rests on such a vast disjunction of different packages of properties of organisms and their environments that we could never define or exhaustively quantify fitness differences in terms of all their different causes. Nevertheless, once we see that fitness supervenes on disjunctions of complex combinations of environmental/organismic relations, the fact that fitness is nothing but such combinations and yet can't be exhaustively defined or explained in their terms (since there are too many disjunctive terms) becomes clear.

Similarly, the fact that we cannot translate functional vocabulary into statements about evolutionary etiologies reflects the fact that a functional kind, such as 'wing' may be the result of a vast disjunction of evolutionary etiologies which differ from one another in all the routes that random variation can take to solving the problem of achieving lift and thrust. But the supervenience of one biological level of description upon another, and ultimately upon the molecular, makes most pressing the general philosophical problems of a) clarifying the ontological status of disjunctive properties—since the supervenience base for any higher level of biological description is necessarily disjunctive owing to the plurality of routes natural selection can take to a single adaptational outcome; and b) establishing the existence of laws in biology below the level of generality of the PNS, or otherwise illuminating the nature of explanation in biology to the extent that it must rely on the description of historical patterns, as opposed to real general laws.

Perhaps the question has occurred to the reader why supervenience begins to appear at the level of the biological and not before. Why is it that the reduction of phenomena at one level of complexity to phenomena at a simpler level proceeds so smoothly from thermodynamics to the kinetic theory of gases, or from the laws of chemical reactions to atomic theory, but not so smoothly from t

he biological to the biochemical, or more generally from function to structure? The answer may by now be clear. Supervenience is a consequence of natural selection. For natural selection selects for effects (whence the “selected effect” account of function). But selection for effects is blind to differences in structure which have the same consequences for persistence and reproduction of the structures themselves. As a result over time, alternative structures will be selected for owing to their similarity in effects. With enough time, the result will be the supervenience of the functional, the biological, on the structural, the physical.

The organization of nature is thus a function of the operation of the PNS on the initial physical conditions that characterized the Earth at the outset of evolution. Accordingly, such regularities as biology uncovers will be historical patterns resulting from the interaction of this law of nature and a set of particular facts about the Earth’s history. Thus, to the extent that these empirical generalizations, whether in molecular or classical genetics, physiology or ecology, have explanatory power, it will derive ultimately from the fact that they are underwritten by the PNS. Notice that this conclusion reconciles the autonomy of the biological from the physical without requiring an account of theory or explanation in which laws are absent. There is a law which does the work required to make biological explanations autonomous from physical ones, and there are important historical patterns we might miss were we to eschew the search for adaptation at levels above the physical. And all this is perfectly compatible with the realization that the physical facts fix all the biological facts.

It will be no surprise that neither classical genetics nor traditional embryology, nor cell physiology, nor any other compartment of biology can be reduced, law by law and theory by theory, to the laws of molecular biology. For there are no laws of either the reduced or the reducing “theory”. There are particular processes—historical patterns—describable at the level of functional biology, and particular processes describable at the level of molecular biology, and each of the former tokens is realized by one of the latter tokens. The explanation for how these patterns came about involves the PNS and the initial conditions under which the particular macromolecules came to be organized by the PNS on the Earth at some time in the evolutionary past.

#### 4. Genic reductionism and the levels of selection.

Both the philosophy of science in general and the philosophy of biology have struggled with the controversy between emergent holists and methodological individualists, not just in the context of the revolution in molecular biology, but also in the context of the explanation of the properties of

f groups—biological and social. Since Durkheim early in the 20<sup>th</sup> century social scientists have argued that there are units of social organization above the individual, which have an existence independent of the individuals that compose them, and whose existence and features can be explained by appeal to Darwinian forces of natural selection. Individualists have rejected these views, and the social scientists among them have often rejected the bearing of Darwinian theory on human affairs. It is ironic therefore that in contemporary biology, it has been popular to hold that Darwinian selection makes impossible the existence of groups and other units of selection higher than the individual, or indeed higher even than the individual gene.

It is in settling the controversy between exponents of the claim that selection operates only at the level of the gene, and those who claim it operates autonomously elsewhere, that the philosophy of biology has perhaps done the most for biology and indeed for the special sciences which rely upon it for their explanations of the existence and features of groups and social institutions that characterize them. We shall show that the result of this discussion in the philosophy of biology is an agreement on a substantive thesis, which makes science safe for some versions of holism: a pluralistic hierarchical approach to selection is the only one that is viable and that monistic genic selectionism is biologically unwarranted. But first some history.

If science starts with experience, then it should not be surprising that biology, and in particular evolutionary biology, should take as its initial problematic the origin of individual organisms of the same order of magnitude of size and complexity as ourselves. It will not be immediately obvious, that where our inquiry starts, with lineages of individual organisms, is not where evolution starts. It must begin with creatures much less highly evolved. Similarly, we cannot take as fundamental to a general evolutionary theory generalizations, models, and equations, such as the Hardy-Weinberg Law, which begin to obtain only once highly evolved organisms appeared bearing mechanisms of meiosis and syngamy (see Beatty [1982]). Similarly we cannot frame a general theory of units upon which selection acts in terms of highly evolved organisms such as ourselves. Developing a theory of evolution has to begin somewhere, but its developers must remember that they seek a theory which accounts for evolution of all biological systems, not merely those of our order of magnitude.

Both for expository purposes, and adventitious ones, Darwin [1859] began with organisms as his paradigmatic objects of selection. It was organisms that varied in subtle ways that could affect their chances of survival and reproduction, and they formed parent-offspring lineages that exhibited heritability. In short, organisms clearly satisfied what has come to be known as Darwin's three co

conditions, (1) there is variation among traits ; (2) some of this variation is heritable, and (3) some variants reproduce more than others. (It was Lewontin 1970 who first boiled down Darwin's theory to these three conditions explicitly).

It is important to note that in a few places Darwin did consider non-organismic levels of selection. On the Origin of Species treats the occurrence of sterile castes among the social insects as the result of selection, not for individual organismal fitness, but fitness of the insect colony, treated as an individual. In The Descent of Man (1871) Darwin speculated that certain human traits, e.g., altruism in battle, evolved by a process of group selection. Identifying the immediate cause of altruism in feelings of sympathy, he wrote,

...this feeling...is one of high importance to all those animals which aid and defend one another, it will have been increased, through natural selection; for those communities, which include the greater number of the most sympathetic members, will flourish best and rear the greatest number of **offspring**. [Descent, pp. 82]

As Darwin clearly understood, this implied that groups varied with respect to quality and/or quantity of altruism shown by their members, that this variation was to some degree heritable, and that this variation resulted in differential survival and reproduction of the groups (as opposed to their component members). But this speculation about human evolution was seen as exceptional by Darwin.

During the time between Darwin and the 1960's a number of ecologists and ethologists argued that certain traits exist for the good of the group, or species. From Peter Kropotkin at the beginning of the 20<sup>th</sup> century, various writers, argued that animal behavior was selected for its group-beneficial consequences. Konrad Lorenz [1950] for example, argued that wolves limit their competition by engaging in ritualized intraspecific non-fatal combat, and thus by avoiding the killing of conspecifics, and they benefit the species. V.C. Wynne-Edwards [1962] argued that individual organisms' **offspring** numbers are regulated in such a way as to maintain the optimal population size of their species. The explanation he offered for such apparently goal-directed activity was selection of those groups in which individual organisms happened to sacrifice their reproductive fitness for that of the group's fitness. Like Darwin's speculations, Lorenz's discussion and Wynne-Edwards' were informal, and without theoretical foundations. This sort of verbal argument made little contact with the growing body of mathematical population genetics. Matters changed in 1966 when the biologist G. C. Williams published Adaptation and Natural Selection. In particular, Williams [1966] subjected Wynne-Edwards' [1962] arguments that animal regulation of population size was due largely to group

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

p selection to a critical examination. He argued that most putative cases of group selection could be explained by individual selection. For example, David Lack [1954] had shown that although many birds do indeed regulate their clutch size (i.e., they lay fewer eggs than they are physiologically capable of laying), this was individually advantageous, since such birds would have more offspring surviving to sexual maturity than birds laying more eggs. More importantly, Williams argued that all selection could be understood at the genic level: the environment really only ever selects for the fitter among competing genes. All other selection is somehow to be cashed in for genic selection. This view was later popularized by Richard Dawkins, The Selfish Gene [1976].

At least one thread of Williams' argument led to a monistic theory of units of selection, namely, all selection is genic selection. This was certainly a departure from the view point with which Darwin began, that selection operated on organisms—biological entities of the same order of magnitude as the human being. But gene selection theory had the virtue of reflecting how natural selection might account for all evolution by focusing on a lineage common to all subjects of selection at least on the Earth. Moreover, it made the relationship between the theory of natural selection and genetics—both population and molecular—particularly straightforward. William's monism was however rather too neat.

In 1970 Lewontin advanced a hierarchical account of the units of selection. He showed how Darwin's basic theory of evolution by natural selection could be put in an abstract and general way that could be applied to account for evolution at a large number of levels of biological organization. So far as Darwin's theory is concerned, any thing will count as a unit of selection that exhibits heritable variance in fitness (Darwin's three conditions above). Organisms clearly satisfy these conditions (not always, but almost always), but so do certain suborganismic entities (e.g., nucleotide sequences, organelles, cell lineages), and, in theory at least, so could certain superorganismic entities (e.g., groups of organisms, perhaps even species). Evolution can therefore proceed on multiple levels. The prospects for conflict between this view and William's claim, that evolution is always and only the natural selection of genes, is evident.

Two suborganismic cases nicely illustrate Lewontin's hierarchical perspective. The first involves understanding the evolution of the large amounts of repetitive DNA sequences that do not code for any protein product, which molecular techniques began to reveal in the 1970's. Two papers published together in Nature in 1980 came to basically the same conclusion, namely that these repetitive sequences evolved not by means of normal organismic selection but rather by means of a withi

non-cellular selection process. In this process nucleotide sequences that could replicate themselves and insert their copies elsewhere in the genome would increase in frequency relative to sequences lacking that ability, whether they made any contribution to the cell, or organism in which they were found. (Doolittle and Sapienza [1980], Orgel and Crick [1980]). Both sets of authors used the term "selfish" taken from Dawkins' book, The Selfish Gene [1976]. Doolittle and Sapienza used the term "selfish genes", Orgel and Crick called these sequences "selfish DNA".

Despite the name, and the evidence they provided for the view that selection can operate directly on gene sequences, these discoveries did not confirm Dawkins' and Williams' monistic view of selection. They actually refuted the view. This is perhaps clearest in Doolittle and Sapienza's [1980] discussion. They note that in order to understand the process that leads to these repetitive sequences one has to first jettison the "phenotypic paradigm", by which they meant the organismic phenotypic paradigm. Asking how such sequences benefit the organismic phenotype leads nowhere. These sequences which do not code for any gene product cannot be beneficial to the organism housing them; indeed they can even be harmful to the organism. The crucial question, and the one which leads to the correct understanding of this selection process is how do such sequences benefit their own lineage by replication and insertion? The answer is obvious. In the environment of the genes, the molecular milieu of the nucleus, such sequences have higher "fitness" than their competitors, i.e., they increase in frequency relative to sequences that do not engage in that behavior.

But, of course, not all selection can be so understood. When camouflage or cryptic coloration increases in a population of moths, that is not because the alleles leading to crypsis out-compete the alternative non-cryptic alleles by the sort of within-cellular process described by Doolittle and Sapienza and Orgel and Crick. Rather, the cryptic colored moths out-compete, and therefore out-reproduce, the more conspicuous moths in their common selective environment. This organismic process leads to increased frequency of the cryptic alleles relative to the non-cryptic. If "selfish DNA" (in the sense of the papers under discussion) were simply an alternative and completely general way of thinking about all selection, as Dawkins (1976) and Williams (1966) seemed to think, then the process described in those papers would not have been a distinct and different process from the sort of selection that increases the frequency of gene sequences that do code for protein products. But this sort of selection requires the selection of lineages of organisms which house these functioning genes. This is what makes the discovery of "selfish DNA" interesting: it constitutes a companion to organismal selection, not a substitute for it.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

The second case of suborganismal selection which in fact requires a hierarchical perspective on evolution is perhaps, more interesting than that presented by “selfish DNA”. As noted, evolution did not begin with organisms such as ourselves. That is, it did not begin with complex, multicellular, diploid, sexual organisms that have a specialized germline that is sequestered early in development. If that is what a biological “individual” is, then individuality is itself an evolved outcome. That is what Leo Buss [1987] argued in an important book. According to Buss, to understand how individuals—i.e. creatures such as ourselves—evolved in the first place, so that there could be selection among us based on such factors as differential camouflage or differential attractiveness to potential mates, we must understand how selection processes could operate on lower levels of organization which had to have existed first in order for items at these levels to be combined into individuals such as us, and to see how such lower level processes could lead to the evolution of higher level entities.

To greatly oversimplify, multicellularity arose when previously independent cells aggregated together in ways that fortuitously enhances the fitness of one another’s lineages. In other words, these aggregations performed some function, previously unavailable to them and advantageous to them. But with multicellularity comes the potential for conflict. If one cell lineage can “selfishly” exploit the work of another to increase its representation in future without expending resources performing some function for the lineage it exploits, that selfish cell lineage will be selected for. The stability of complex multicellular organisms depends on somehow policing this sort of conflict. One way of doing this, a way that has been settled on by a numerical minority of organisms on the Earth is first, to go through life cycles with single-cell bottlenecks so that all cells derived from the single initial cell start out genetically identical, and second, to sequester the germ cells early in development so that they are insulated from any competition among cell lineages in the organism’s later life. In this way, even if mutation introduces genetic differences among the developing cells, competition among genetically different cell lineages would have no evolutionary consequences since the germ cells are insulated from this competition. Imagine your right arm in competition with your left and assume they derive from genetically differentiated cell lineages. Given the sequestration of your germ cells, this competition could have no effect on the genetic composition of your offspring.

Recent work on major evolutionary transitions—transitions from one level of organization (e.g., the cell) to another (e.g., multicellular individuals)—reinforces the point we are making here (see Maynard Smith and Szathmary [1995]). That is, a general theory of evolution by natural selection must allow for evolution to be hierarchical. The virtual consensus on the need for hierarchy raises

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

the question: A hierarchy of what? Lewontin's [1970] seminal work suggested a hierarchy of "units" of selection, where a biological item is a unit just in case it satisfied Darwin's three conditions. It is here that problems arise on the intersection of philosophy and biology. For Lewontin's claim was too strong, as philosophers began to uncover.

Wimsatt [1980, 1981] uncovered the problem in his development of Lewontin's analysis. To repeat, Lewontin had argued that any entity that had heritable variation in fitness, i.e., any entity that satisfied Darwin's three conditions, would be a unit of selection. Wimsatt showed that the satisfaction of Darwin's three conditions was necessary, but not sufficient for the existence of a unit of selection, because an entity that was itself composed of units of selection could satisfy the conditions without itself being a unit of selection. For instance, a group of plants might have a high average height compared to other nearby groups of that species, and might thereby have greater resistance to herbivory than the other shorter groups (let us suppose the herbivores themselves are short). This leads to greater reproductive success of the taller groups relative to shorter groups. Furthermore, suppose that new groups are formed from existing groups in a way that leads to heritability of average height among groups. Here the groups of plants have heritable variation in fitness. But is this really group selection? From the description given so far we cannot say. If individual plant fitnesses vary in a way that is not dependent on group context, i.e., if all that matters to selection of a plant's lineage is how tall the plant is, and the average height of its group matters not at all once its height is fixed, then the groups of plants do not qualify as units of selection. If, on the other hand, group membership does matter with respect to individual plant fitness, then the groups do qualify as units of selection. (In Wimsatt's terms, if the variance in individual fitness is context-dependent, then there must be a higher level unit of selection—in this case the group.) Sober [1984] gave a similar analysis of units of selection, and Lloyd [1988] further refined Wimsatt's analysis.

These criteria for the existence of distinct units of selection inspired by Lewontin [1970], (and also [1974]) seemed basically right with respect to whether or not some superorganismic (e.g., group) level of selection exists. But when it came to addressing the question of whether a case of selection was organismic or suborganismic they get matters quite wrong. Lewontin [1974] and Wimsatt [1980, 1981] advert to certain facts about genetics, especially epistasis and gene linkage, in arguing against genic selectionism. But the conclusion supported by these arguments was that some larger genetic unit, a length of chromosome, perhaps, in the limit the entire genome, was the unit of selection in garden variety cases of selection. However, the proper contrast is not between the genotype a

nd the organism, but between at least three different units of selection: our normal cases, say of cryptically colored moths, the suborganismic cases, and the superorganismic cases, that is between organismic selection versus group selection versus genic selection (as in the cases of Orgel and Crick and Doolittle and Sapienza discussed above).

Though this conclusion is a step in the right direction, it fails to distinguish two separate questions and two separate processes. Hull [1981], Dawkins [1982] and Brandon [1982] independently recognized that the prior discussions of units of selection had been flawed by the conflation of two distinguishable questions. In Hull's terminology, we must distinguish the question of what entities are interacting directly with their environment in a way that leads to differential reproduction (interactors) from the question of what entities are reproducing their structure directly and accurately (replicators). (There are substantive and terminological differences among these three authors. The substantive differences are minor compared to the large area of mutual agreement and will not be discussed here, but see Brandon [1990], chapter 3 for further discussion). Terminologically Dawkins made the distinction between replicators and vehicles, while Brandon distinguished between units and levels of selection. Hull's terminology seems to be the most commonly adopted and will be used here.) Thus it turns out that there is not a single hierarchy with respect to selection, but a dual hierarchy, one of interactors and one of replicators (see Brandon [1988] and [1990]).

The interactor question clearly is a question of causality. What makes a case of selection organismic, as opposed to genic or group, is that the differential reproduction of organisms is caused by differences in organismic phenotypes. For example, differences in cryptic coloration cause differences in survival and reproduction in some moth populations. Likewise, what makes a case of group selection is that differences in group phenotype (e.g., amount of altruism expressed within groups) causes differences in group survival and reproduction. Brandon [1982] used the probabilistic notion of screening-off to explicate this basic idea. Sober [1984] employed a Pareto-style analysis of causation and not surprisingly there are differences between the two accounts (see Sober [1992] and Brandon et al. [1994]). But these differences should not blind us to the consensus that has been reached. There is a hierarchy of potential interactors ranging from bare RNA sequences up to species and even clades. Further we have an abundance of evidence of organismic level selection in addition to sufficient evidence of several different sorts of suborganismic interactors. A few cases of group selection have been documented. At this point we think the interactor question is sufficiently well defined so that the difficult job of answering the biological questions of how frequent are certain n

on organismic levels of selection can now proceed. And it is proceeding (see e.g., McCauley [1994 and Stevens et al. [1995]). On the empirical issue of how frequent, or how important, are processes of superorganismic selection philosophers may remain agnostic

Another way of thinking about the interactor question is to think about it from the point of view of evolutionary adaptation. An adaptation is a trait that has evolved because of the selective advantage it confers on its bearers (Brandon [1978, Lewontin [1978). In that sense, ‘adaptation’ is an explicitly causal-historical term (just as calling a mountain volcanic is making a claim about its causal history). In contrast, to call a trait adaptive is to simply make a claim about its current utility. Darwin’s theory of evolution by natural selection had as one of its primary goals explaining the existence of adaptive traits. Most adaptive traits exist because they are adaptations. One of the great advantages of a hierarchical theory of natural selection is that it allows us to distinguish adaptations at different levels of biological organization. What is good at one level may often not be good at another level. As discussed earlier, what is good for one cell lineage in a multicellular organism may not be good for the organism as a whole (think of cancer). Altruism, as Darwin clearly saw, is of no benefit, indeed will usually be of harm, to the individual organism, but can be of benefit to the group as a whole.

Dawkins ([1976] and [1982] proposed that active germ-line replicators were the ultimate beneficiaries of all adaptations. Brandon ([1985, [1990] argued that this view blinds us to important differences in the immediate targets of selection. If all adaptations are for the good of active germ-line replicators, there is no way to distinguish genuine cases of genic adaptations (as in the cases of Dolittle and Sapienza [1980, and Orgel and Crick [1980], from organismic adaptations, from group adaptations (possibly, human altruism). Moreover, selection may move in one direction at one level and in the opposite level at another. Thus, altruism may be adaptive for the group and maladaptive for the individual. Given the potential conflict that exists among levels we must make such distinctions in order to understand the evolution of adaptations. Genic selectionism achieves at most an empty generality at the cost of explanatory vacuity. Only the hierarchical approach can offer an explanatory theory of adaptations.

In their review of philosophical work on units of selection Sober and Wilson state that objects “at level  $\times$  are units of selection in the evolution of trait  $T$  [if and only if] one of the factors that influenced  $T$ ’s evolution was that  $T$  conferred a benefit on objects at level  $\times$ ” (’94 ([1994], p. 536).

That formulation fits very nicely the hierarchical approach to interactors.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

What about replicators? Once the question about the level or levels of causal interaction between interactor and environment is settled for a particular case, and the mode of reproduction of these interactors is understood, then we can determine what objects are replicators in that case. For instance, when the level of selection is organismic, i.e., the interactors are organisms, then the replicators are determined by the mode of reproduction, in this case, sexual or asexual. In sexual reproduction, the genomes of the two parents are both broken up by meiosis, resulting in haploid sex cells. Two such cells, one from each parent, are then joined by syngamy resulting in a new diploid cell, genetically different from both parents. That is why the genome as a whole is not a replicator when reproduction is sexual. Instead, a smaller genetic unit, what both Dawkins [1976] and Williams [1966] define as a gene, is the replicator in these cases. But when reproduction is asexual the genome as a whole can be considered a replicator. Similar distinctions would have to be made in cases where groups are interactors—depending on how they reproduce different entities will count as replicators. The point is that corresponding to a hierarchy of potential interactors there is a (not so neat) hierarchy of replicators (see Brandon [1990, chapter 3 for fuller discussion). Thus instead of a single hierarchy of “units” of selection (*sensu* Lewontin), we have a dual hierarchy of interactors and replicators. (See also Griesemer, *in press*).

##### 5. Conclusion: the past and the future of philosophy of biology.

Philosophy of biology has made remarkable progress over the last twenty-five years. We now understand, in rough outlines at least, the special nature of evolutionary theory and the evolutionary process. Natural selection is what differentiates biological evolution from other sorts of physical processes. And natural selection makes the relation between the biological and the physical so complex as to suggest some sort of nonreductive physicalism as a philosophy of biology. Natural Selection is also what underwrites a purposive notion of function, one that squares naturalism with a sort of Darwinian teleology. And the schematic nature of the Principle of Natural Selection best explains both the great generality of the theory of evolution by natural selection and its particular empirical content when applied to particular populations within particular environments. Finally, once clear on the basic structure of the theory we can see how to expand its hierarchical operation, as the evidence about life on this planet demands. These accomplishments are substantial and, although there is room for fine-tuning, we do not expect similar progress in these areas over the next twenty-five years. What then of the future?

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

Philosophers of biology can expect the continuing revolution in biology to generate new problems on the intersection of its concerns and biology's. Most obviously, there are huge questions in developmental biology, but probably in other areas as well. Developmental biology promises to fill in certain areas which evolutionary biologists have treated as so many black boxes. What is the nature of variation? How do point mutations turn into organismic variations? On the path from archaea to metazoans, which evolutionary transitions are most likely, which least likely? What are the "natural" traits that evolution operates on? Will deeper knowledge of development reveal new principles that must be added to a complete evolutionary theory? Will it "decenter" molecular genetics? The questions just asked all relate to the impact of deeper understanding of developmental biology on evolutionary biology. But, of course, developmental biology is worthy of study in its own right. And it raises questions about developmental systems, about whether viewing the genome as preformational versus informational provides the right understanding of the genome, about whether germ-plasm bottlenecks are required for the evolution of complex adaptations. Some of these are questions on which philosophers have already had something to say. For example, Griffiths and Gray, [1994] have argued that folding development into evolution completely undermines the traditional genetic/non-genetic, informational/non-informational taxonomy of molecular biology. Rosenberg [2000a] argues on the contrary that discoveries in molecular developmental biology vindicate a reductionist evolutionary methodology more completely than traditional arguments from molecular genetics. These are first words, not last words however.

Similarly, ecology seems philosophically rich, yet relatively unexplored by philosophers of biology. On the one hand, there are the same sort of conceptual and methodological issues that have kept philosophers of evolutionary biology occupied for many years. For example, what, if anything, unifies ecology? What is the relation between ecology and evolutionary biology? What is the relation among different levels of ecological analysis, ranging from the community level to the whole biosphere? On the other hand, answers to basic questions in the philosophy of ecology have almost immediate ramifications for ethical and policy issues. For instance, what is diversity? How can it be quantified? Why should we care to preserve it? These questions merely scratches the surface. There is much room here for useful philosophical work.

In this review we have not discussed the philosophical work that has been produced relating to questions in systematics (most especially the "species problem"). But systematics has undergone a thorough revolution within the last twenty-five years. A revolution that has involved method (cla

distic); technology (high-speed computers); and new data sources—both protein and DNA sequence data for example (see Sober [1988]). It is hard to see how these revolutionary developments could fail to be of epistemological and methodological significance. And this revolution in systematics spills over into the human genome project and its aftermath. How will we deal with the massive, purely numerical data that has resulted from the sequencing of "the" human genome? Automated methods of hypothesis generation will surely be developed, just as they have been successfully developed in systematics. What does this sort of algorithmic data mining tell us about science? For one thing, it suggests that science isn't all hypothetico-deductive, as many textbook treatments of science would have it.

We have paused to reflect on the accomplishments, and remaining problems, of recent philosophy of biology. A similar review article, written twenty-five years from now will, we predict, look much different.

Robert Brandon  
 Alex Rosenberg  
 Department of Philosophy  
 Duke University  
 Box 90743  
 Durham, N.C., 27708

#### References

- Amundsen, R., and Lauder, G. [1998]: "Function without Purpose" in Allen, C., Bekoff, M., and Lauder, G., (eds), *Nature's Purposes*, Cambridge, MIT Press.
- Cummins, R [1975]: "Functional Analysis", *Journal of Philosophy*, **72**, pp 741-65.
- Beatty, J.: 1980. 'What's wrong with the received view of evolutionary theory?', in P. D. Asquith and R. N. Giere (eds.), *PSA 1980*, Volume 2, Philosophy of Science Association, East Lansing MI, pp. 397-426.
- Beatty, J. [1983]: "Insights and Oversights of Molecular Genetics" in Asquith, P., and Nickels, T. (eds), *PSA 1982*, v. 2, East Lansing, Mi., Philosophy of Science Association.
- Beatty, J. [1992]: in E. Fox-Keler, and L. Lloyd (eds), *Keywords in Evolution*, Cambridge, Ma., Harvard University Press.
- Boch, W.S., and Wahlert, G., [1965]: "Adaptation and the Form-Function Complex", *Evolution*, **19**, pp. 269-299.
- Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

- Brandon, R. N.: [1978]: 'Adaptation and evolutionary theory', *Studies in History and Philosophy of Science* **9**, 181-206.
- Brandon, R., [1982]: 'The levels of selection', in P. Asquith and T. Nickles (eds.), *PSA 1982, Vol. 1*, Philosophy of Science Association, East Lansing, MI, pp. 315-323.
- Brandon, R. [1985]: 'Adaptation explanations: Are adaptations for the good of replicators or interactors?', in B. Weber and D. Depew (eds.), *Evolution at a Crossroads: The New Biology and the New Philosophy of Science*, The MIT Press, Cambridge, MA, pp. 81-96.
- Brandon, R. [1988]: 'The levels of selection: A hierarchy of interactors', in H. Plotkin (ed.), *The Role of Behavior in Evolution*, The MIT Press, Cambridge, MA, pp. 51-71.
- Brandon, R., [1990]: *Adaptation and Environment*, Princeton University Press, Princeton.
- Brandon, R. N., Antonovics, J., Burian, R., Carson, S., Cooper, G., Davies, P. S., Horvath, C., Mishler, B. D., Richardson, R. C., Smith, K. and Thrall, P. [1994]: "Discussion: Sober on Brandon on screening-off and the levels of selection", *Philosophy of Science* **61**, 475-486.
- Buss, L. W.: [1987]: *The Evolution of Individuality*, Princeton University Press, Princeton.
- Campbell, D.T. [1974]: "Evolutionary epistemology", in *Methodology and Epistemology for Social Science* Chicago, University of Chicago Press.
- Darwin, C. [1859]: *On the Origin of Species*, John Murray, London.
- Darwin, C. [1871]: *The Descent of Man*, John Murray, London.
- Dawkins, R. [1976]: *The Selfish Gene*, Oxford University Press, Oxford.
- Dawkins, R. [1982]: 'Replicators and vehicles', in King's College Sociobiology Group (eds.), *Current Problems in Sociobiology*, Cambridge University Press, Cambridge, pp. 45-64.
- Dobzhansky, T. [1972]: "Nothing makes sense except against the background of evolution", *American Biology Teacher*, **35**, pp. 125-129.
- Doolittle, W. F. and Sapienza, C. [1980]: 'Selfish genes, the phenotype paradigm and genome evolution', *Nature* **284**, 601-603.
- Dreisch, W. [1914]: *History and Theory of Vitalism*, London, Macmillan.
- Ghiselin, M [1969]: *Triumph of The Darwinian Method*, Berkeley, University of California Press.
- Griesemer, J.: in press a, 'Development, culture and the units of inheritance', *Philosophy of Science (Proceedings)*.
- Griesemer, J. in press b, 'Reproduction and the reduction of genetics', in P. Beurton, R. Falk and H. J. Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)

. Rheinberger (eds.), *The Concept of the Gene in Development and Evolution*, Cambridge University Press, Cambridge.

Griffith, P, and Grey, R. [1994]: “Development systems Theory and Evolutionary Explanation”, *Journal of Philosophy*, **91**, pp. 277–304.

Gillsespie, [1975]: “Natural Selection for Within-Generation Variance in **Offspring** Numbers”, *Genetics*, **76**, pp. 601–606, **81**, pp. 403–413.

Hull, D. [1981]: ‘Units of evolution: a metaphysical essay’, in U. L. Jensen and R. Harrington (eds.), *The Philosophy of Evolution*, Harvester Press, Brighton, pp. 23–44.

Keynes, J.M. [1921]: *A Treatise on Probability*, London, Macmillan.

Kitcher, P. [1984]: “1953 and all that: A Tale of Two Sciences”, *Philosophical Review*,

Kitcher, P. [1999]: “The Hegemony of Molecular Biology” *Biology and Philosophy*

Lack, D. [1954]: ‘The evolution of reproductive rates’, in J. S. Huxley, A. C. Hardy and E. B. Ford (eds.), *Evolution as a Process*, Allen and Unwin, London, pp. 143–156.

Lennox, J. [1992]: “Teleology”, in E. Fox-Keller, and L. Lloyd, (eds), *Keywords in Evolution*, Cambridge, Ma., Harvard University Press.

Lewontin, R. C. [1970]: ‘The units of selection’, *Annual Review of Ecology and Systematics* **1**, 1–18.

Lewontin, R.C. [1974]: *The Genetic Basis of Evolutionary Change*, Columbia University Press, New York.

Lewontin, R.C.[1978:]: ‘Adaptation’, *Scientific American* **239**, 156–169.

Lloyd, E. A.. [1988]: *The Structure and Confirmation of Evolutionary Theory*, Greenwood Press, New York.

Lorenz, K. [1950]: ‘Ganzheit und Teil in der tierischen und menschlichen Gemeinschaft’, *Studium Generale* **3/9**.

Mayr, E. [1961]: “Cause and Effect in Biology”, *Science*, **134**, pp. 1501–6.

McCauley, D. E. [1994]:. Intrademic group selection imposed by a parasitoid–host interaction. *American Naturalist* **144**: 1–13.

Maynard Smith, J. and Szathmáry, E. [ 995]: *The Major Transitions in Evolution*, W. H. Freeman, Oxford.

Michod, R.E. [1985]: “On Adaptedness and Fitness and their Role in Evolutionary Theory”, *Journal of the History of Biology*, **19**, pp. 289–302.

*Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)*

- Mills, S., and Beatty, J., [1979]: “The propensity interpretation of fitness”, *Philosophy of Science*, **46**, pp. 263–286.
- Nagel, E. [1961]: *Structure of Science*, New York, Harcourt Brace and World.
- Orgel, L. E. and Crick, F. H. C. [1980]: ‘Selfish DNA: The ultimate parasite’, *Nature* **284**, 604–607.
- Rosenberg, A. [1978]: “The Supervenience of Biological Concepts”, *Philosophy of Science*, **45**, pp. 368–386.
- Rosenberg, A. [1985]: *The Structure of Biological Science*, Cambridge, Cambridge University Press.
- Rosenberg, A. [2000], “Reductionism Redux”, in *Darwinism in Philosophy, Social Science and Policy*, Cambridge, Cambridge University Press.
- Rosenberg, A. [2000a], “Reductionism in a Historical Science”, *Philosophy of Science*, **68**, forthcoming.
- Ruse, M. [1973]: *The Philosophy of Biology*, London Hutchinson.
- Simpson, G. C. [1949]: *The Meaning of Evolution*, New Haven, Yale University Press.
- Stevens, L., Goodnight, C. J., and Kalisz, S. [1995]: Multilevel selection in natural populations of *Impatiens capensis*. *American Naturalist* **145**: 513–526.
- Sober, E. [1984]: *The Nature of Selection*, The MIT Press, Cambridge, MA.
- Sober, E. [1992]: ‘Screening-off and the units of selection’, *Philosophy of Science* **59**, 142–152.
- Sober, E. [1993]: *The Philosophy of Biology*, Boulder, Co, Westview Press.
- Sober, E. and Wilson, D. S. [1994]: ‘A critical review of philosophical work on the units of selection problem’, *Philosophy of Science* **61**, 534–555.
- Sommerhoff, W. [1950], *Analytical Biology*, NY, Oxford 1950.
- Szathmari, E. and Maynard Smith, J.: [1997] ‘From replicators to reproducers: the first major transitions leading to life’, *Journal of Theoretical Biology* **187**, 555–571.
- Thompson, P. [1988]: *The Structure of Evolutionary Theory*, Albany, SUNY Press.
- Williams, G. C. [1966]: *Adaptation and Natural Selection*, Princeton University Press, Princeton.
- Wimsatt, W. C.: [1980]: ‘Reductionistic research strategies and their biases in the units of selection controversy’, in T. Nickles (ed.), *Scientific Discovery*, Volume II, Historical and Scientific Case Studies. Do not quote without approval of author [alexrose@duke.edu](mailto:alexrose@duke.edu)

udies, Reidel, Dordrecht, pp. 213–259.

Wimsatt, W., [1981]: ‘The units of selection and the structure of the multi-level genome’, in P. D. Asquith and R. N. Giere (eds.), PSA 1980, Vol. 2, Philosophy of Science Association, East Lansing, MI, pp. 122–183.

Wright, L [1976]: Teleological Explanation, Berkeley, UC Press, 1976].

Wynn-Edwards, V. C. [1962]: Animal Dispersion in Relation to Social Behavior, Oliver and Boyd, Edinburgh and London.