

Objectifying values in science: A case study

Mark A. Bedau

Reed College, 3203 SE Woodstock Blvd., Portland OR 97202

503-517-7337

<http://www.reed.edu/~mab>

mab@reed.edu

Background to an objectification of values

There are at least two different ways in which values and science can be connected. One is through the evaluation of science, and the other is through the scientific investigation of values. The evaluation of science is a non-scientific, political or ethical investigation of the practices of science. Various proposed and actual scientific practices call out for social and ethical evaluation. A few that have received recent attention are the human genome project, intelligence testing, and encryption algorithms. Such evaluations of science contrasts sharply with what I call “the science of values.” This is not one science or even one unified nexus of scientific activities but a loosely defined grab bag containing all scientific investigations of matters involving values.

One part of the science of values concerns what individuals or groups value or take an interest in—these are values considered from a first-person point of view. The values can concern anything including morality, aesthetic matters, religion, politics, lifestyle, livelihood, etc. The science of first-person values includes such things as psychological studies of the values of individual people, sociological studies of the values of social groups, and anthropological comparisons of the values of different cultures.

Another part of the science of values concerns what is good for, or promotes the interests of, some individual or group from an external, third-person point of view. The subject whose interests are being studied might or might not internalize the values used in the external evaluation. Examples of the science of third-person values include studies of the value of a college education, of regular visits to the dentist, or of growing and eating organic food. They also include biological studies of what kinds of traits help creatures to survive, reproduce, and generally flourish. This latter example is directly connected to the story I will relate here. This story concerns my own participation in the science of values over the past decade, which grew out of a value-centered theory of biological teleology I developed fifteen years ago.

My story starts in 1991 when I was invited to present my theory of teleology to the newly founded artificial life research groups at Los Alamos

National Laboratory (LANL) and the Santa Fe Institute (SFI). Artificial life is an interdisciplinary endeavor that studies life and life-like processes by simulating them or synthesizing them. Much of this work consists of computer models of processes like the self-organization of simple abstract metabolisms or the evolutionary dynamics of populations of simple self-reproducing automata. Artificial life aims to understand the essential properties of the fundamental processes at work in any possible living system. As Chris Langton once put it, its goal is to understand not “life-as-we-know-it” but “life-as-it-could-be.” Pursuing this goal requires having a general and broad grasp of what life is and could be, so the LANL-SFI artificial life group created a seminar on the nature of life. But the group was unable to formulate an adequate definition. Disappointed, the best they could produce was a list of characteristic hallmarks of life (see Farmer and Belin 1992).

Teleology in one form or another is often considered one of the hallmarks of life (see, e.g., Monod 1971, Mayr 1982), but the notion of teleology is no more self evident than the notion of life itself. So, knowing about my work, the LANL-SFI artificial life group invited me to present to them a philosopher’s perspective on teleology. I knew that the artificial life group did not consider itself to understand a theory fully unless it could see how to implement it in a computer model, so I augmented my presentation with a discussion of how to operationalize the key elements of my theory. My theory of teleology concerns traits that are explained by their value, so operationalizing the theory consisted in figuring out how to determine objectively and impartially when a trait’s value or usefulness explains its continued existence. The “objectifying value” of my title refers to this kind of operationalization.

After my lecture, Norman Packard came up and said that he thought it would be easy to objectify teleology in his computer model of sensory-motor evolution. We worked out the details that night and had our first results the next day (see Bedau and Packard 1992). That was an eye-opening episode that convinced me of the usefulness of operationalizing philosophical theories, whenever possible. This paper describes how I objectified value in biology, illustrates the method in a simple evolutionary system consisting of self-replicating computer programs, and then explains two fruits of this exercise. One concerns Gould and Lewontin’s challenge to adaptationism. The other concerns comparing evolutionary creativity in biological and cultural evolution.

Objectifying teleological explanations in biology

My example of objectifying value in science consists of objectifying the value in a certain kind of biological teleology, specifically, the teleology involved in adaptationist explanations.¹ Traits or behavior that can be explained by reference to the utility of their effects are teleological (telic, for the sake of an end), by my lights.² In ordinary parlance, telic explanations are offered for a

wide variety of things. These include such things as the actions of conscious human agents, and the structure and behavior of artifacts designed and used by people. They also include the behavior and structure of biological organisms, as well as certain lower-level components such as genes and also certain higher-level groups such as populations and species. All of these can have telic explanations, and in each case the beneficial effect brought about by the explanandum is an essential part of the explanation. Functionality or adaptiveness is sometimes confused with teleology; the two are related but different. Functional or adaptive behavior is just behavior that is beneficial, that “serves a purpose,” regardless of how it comes about. Telic behavior, on the other hand, is not merely beneficial, does not merely serve a purpose. It occurs specifically because it is beneficial, because it serves a purpose. Telic behavior cannot occur merely accidentally or for some reason unconnected with its utility. Analogous considerations distinguish merely functional or adaptive traits from telic traits.

My concern in the present context is how this framework applies to biological teleology, in particular. A range of behaviors or traits of a given organism at a given time are more or less adaptive. If an organism contains a favorable mutation, the new behavior or trait caused by this mutation might immediately be adaptive or beneficial. But that behavior or trait will not be telic until its utility becomes a causal factor in its continual production. This can happen if its behavior persists through a lineage because of its utility.

What I am describing, of course, is the process by which natural selection produces adaptations. An adaptation is a trait (possibly a kind of behavior) that is produced by the process of natural selection for that trait.³ For example, the whale’s fins are an adaptation for swimming. The trait persists due to natural selection because of its beneficial effects for swimming; this benefit explains why it is a product of natural selection. Although traits of individual organisms are the paradigm example of adaptations, we can apply the notion to higher level entities by averaging over traits and organisms. In particular, below I will talk of genotypes (the complete set of traits in an organism) as adaptations. A genotype is an adaptation if it persists through the action of natural selection, that is, if on average the individuals with that genotype have been selected for their possession of that genotype, i.e., if the traits in that genotype are adaptations.

The crux of my method for objectifying biological teleology is to observe the extent to which items resist selection pressures, for resistance to selection is evidence of adaptation. Since an item is subjected to selection pressure only when it is active or expressed, I call this evolutionary “activity” information.⁴ Simple bookkeeping collects an historical record of items’ activity – the extent to which items have been subjected to selection pressure, i.e., the extent to which their adaptive value has been tested. The bookkeeping increments an item’s current activity as long as it persists, yielding its cumulative activity. If the item (e.g., gene) is inherited during reproduction, its cumulative activity continues to be incremented by the child’s current activity. In this way our bookkeeping records an item’s cumulative activity over its

entire history in the lineage. Cumulative activity sums the extent to which an item has been tested by selection over its evolutionary history.

Every time an item is exposed to natural selection, selection can provide feedback about its adaptive value. Obviously, an item will not continue to be tested by natural selection unless it has passed previous tests. So, the amount that an item has been tested reflects how successfully it has passed the tests. If a sufficiently well tested item persists and spreads through the population, we have positive evidence that it is persisting because of its adaptive value. That is, we have positive evidence that it is an adaptation, that it is telic.

But natural selection is not instantaneous. Repeated trials might be needed to drive out maladaptive items. So exposure to some selection is no proof of being an adaptation. Thus nonadaptive items will generate some “noise” in evolutionary activity data. To gauge resistance to selection we must filter out this nonadaptive noise. We can do so if we first measure how activity will accrue to items persisting due just to nonadaptive factors like random drift or architectural necessity. A general way to measure the expected evolutionary activity of nonadaptive items is to construct a neutral model of the target system: a system that is similar to the target in all relevant respects except that none of the items in it has any adaptive significance. (I give concrete examples below.) The accumulated activity in neutral models provides a no-adaptation null hypothesis for the target system that can be used to screen off nonadaptive noise. If we observe significantly more evolutionary activity in the target system than in its neutral shadow, we know that this “excess” activity cannot be attributed to nonadaptive factors. It must be the result of natural selection, so the items must be adaptations.⁵

An illustration of evolutionary activity of genotypes

I will illustrate the evolutionary activity test for adaptations in Evita, a simple artificial evolving system that “lives” in a computer.⁶ Somewhat analogous to a population of self-replicating strings of biochemical RNA, Evita consists of a population of self-replicating strings of customized assembly language code and residing in a two-dimensional grid of virtual computer memory. When Evita runs, it doles out CPU time to all programs residing in memory. The system is initialized with a single self-replicating program. CPU time causes this ancestral program to execute, and it copies each of its instructions into a neighboring spot on the grid, thereby producing a new copy of the program—its “offspring.” Then this offspring and its parent both are allocated CPU time and start executing, and each makes another copy of itself, creating still more offspring. This process repeats indefinitely. When space in computer memory runs low and offspring cannot find unoccupied neighboring grid locations, the older neighbors are randomly selected and “killed” and the offspring move to the vacated space. Innovations enter the system through point mutations. When a mutation strikes an instruction in a program, the instruction is replaced by another instruction chosen at random, so mutation is

continually spawning new kinds of programs.⁷ Many are maladaptive but some reproduce more quickly than their neighbors, and these tend to spread through the population, causing the population of strings to evolve over time.

Evita is explicitly designed so that the programs interact only by competing for execution time and space in memory. There is no fixed generation time for Evita programs; some replicate faster than others. On average, those that reproduce faster will supplant their reproducing neighbors. Most significant adaptive events in Evita are changes in reproduction rate, so for present purposes a genotype's fitness can be equated with its reproduction rate. Evita has a clear distinction between genotype and phenotype. A given genotype is simply a string of computer code. If two programs differ in even one instruction they have different genotypes. But two genotypes might produce exactly the same behavior – the same phenotype. If a program includes instructions that never execute, these instructions can mutate freely without affecting the operation of the program. Thus multiple genotypes – without phenotype distinction and so with exactly the same fitness – may then evolve through random genetic drift.

To gather evolutionary activity data in Evita two issues must be settled. First, one must decide which kind of item to observe for adaptations. We will observe whole genotypes. Second, one must operationalize the idea of a genotype's being tested by natural selection. A plausible measure of this is concentration in the population. The greater the genotype's concentration, the more feedback that selection provides about how well adapted it is. A genotype's cumulative evolutionary activity, then, is just the sum of its concentration over time.

In order to discern how much of Evita's genotype activity can be attributed to the genotypes' adaptive significance, we create a "neutral shadow" of it (recall the discussion above). The neutral shadow is a population of nominal "programs" with nominal "genotypes" existing at grid locations, reproducing and dieing. These are not genuine programs with genuine genotypes; they contain no actual instructions. Their only properties are their location on the grid, their time of birth, the sequence of reproduction events (if any) they go through, and their time of death.

Each target Evita run has a corresponding neutral shadow.⁸ Certain events in the target cause corresponding events in the shadow, but events in a shadow never affect the target (hence, the 'shadow' terminology). The frequency of mutation events in the shadow is copied from the Evita target. Whenever a mutation strikes a shadow "program" it is assigned a new "genotype." The timing and number of birth and death events in the neutral shadow is also patterned exactly after the target. Shadow children inherit their parent's "genotype" unless there is a mutation, in which case the shadow child is assigned a new "genotype." The key difference is that, while natural selection typically affect which target program reproduces, random selection determines which shadow "program" reproduces. So shadow genotypes have no adaptive significance whatsoever; their features like longevity and concentration – and hence their evolutionary activity – cannot be attributed to their adaptive

significance. At the same time, by precisely shadowing the births, deaths, and mutations in the target, the neutral shadow shows us the expected evolutionary activity of a genotype in a system exactly like Evita except for being devoid of natural selection. The neutral shadow defines a null hypothesis for the expected evolutionary activity of genotypes affected by only non-adaptive factors such as chance (e.g., random genetic drift) or necessity (e.g., the system's underlying architecture).

Figure 1 about here

Evita's evolutionary graphs depict the history of the genotypes' activity in a given Evita run.⁹ Whenever one genotype drives another to extinction by competitive exclusion, a new wave arises as an earlier one dies out. Multiple waves coexist in the graph when multiple genotypes coexist in the population, and genotypic interactions that affect genotype concentrations are visible as changes in the slopes of waves. The point to appreciate is that the big waves correspond to main adaptations among the genotypes. We can see this clearly in Figure 1 by comparing a typical Evita evolutionary activity graph (above) with an activity graph of its neutral shadow (below). These graphs are strikingly different.¹⁰ Leaving aside the ancestral wave, the highest waves in the Evita are orders of magnitude higher than those in the neutral analogue. This is clear evidence of how the size of a genotype's evolutionary activity waves in Evita reflects the genotype's adaptive significance. In the Evita target, at each time one or a few genotypes enjoys a special adaptive advantage over their peers, and this is reflected by their correspondingly huge waves. The change in dominant waves reflects a new adaptation out competing the prior dominant adaptations. In the neutral analogue, by contrast, a genotype's concentration reflects only dumb luck, so no genotype activity waves rise significantly above their peers.

Figure 2 about here

Figure 2 shows more detail of the evolutionary activity during the beginning of the Evita run in Figure 1, with the average population fitness graphed below. The activity graph is dominated by five main waves, the first corresponds to the ancestral genotype and the subsequent waves correspond to subsequent adaptations.¹¹ Miscellaneous low-activity genotypes that never claim a substantial following in the population are barely visible along the bottom of the activity plot. Comparing the origin of the waves with the rises in average population fitness shows that the significant new waves usually correspond to the origin of a higher fitness genotype. Detailed analysis of the specific program that makes up the genotypes with high activity, we see that the major adaptive events consist of shortening a genotype's length or copy loop.

Figure 3 about here

The moral, again, is that significant evolutionary activity waves are significant adaptations. They correspond to genotypes that are persisting and spreading through the population because of their relative adaptive value. Natural selection is promoting them because of their relative reproduction rate; they flourish because of selection for this, so they are adaptations. The evidence for the moral has three parts. First, new significant waves coincide with significant jumps in average population fitness. This shows that the new genotype spreading through the population and making the new wave is an adaptive advantage over its predecessors. Second, microanalysis of the genotypes in the new waves reveals the genetic novelties that create their adaptive advantage. Third, in a neutral model in which chance and architectural necessity are allowed full reign and natural selection is debarred by fiat, no genotypes make significant waves. So, the major evolutionary activity waves in Evita could be produced only by continual natural selection of those genotypes, and natural selection of the genotypes must be due to selection for their adaptive value.

Neutral variant genotypes are an exception to this moral, but they prove the rule. Notice that the second fitness jump in Figure 2 corresponds to dense cloud of activity waves. Figure 3 is a blowup of these waves. The genotypes in this cloud differ from each other only by mutations at an unexpressed locus, so they all use exactly the same algorithm. They are neutral variants of one another – different genotypes with exactly the same phenotype. So the neutral variants are one and the same phenotypic adaptation. Each genotypic instances of the phenotype is an adaptation because it is persisting due to its adaptive value.

The foregoing is an especially simple example of the evolutionary activity test for adaptations, but it is just one example. The test applies to more than just Evita and the like. It can be applied to natural systems as well as artificial systems. For example, later in this paper we will apply it to the evolution of technology as reflected in patent records, and it has also been applied to evolution in the Phanerozoic age as reflected in the fossil record (Bedau et al. 1997, 1998). It would be straightforward to apply it to various biological populations, if the relevant raw data could be gathered. For example, Richard Lenski has been directing a study of long-term evolution in Escherichia coli (Lenski and Travisano 1994, Cooper and Lenski 2000). Lenski's E. coli populations have been evolving for over 20,000 generations. If we had a record of the concentration of each genotype in these populations every hundred generation or so, then we could plot the evolutionary activity of these genotypes exactly as we did for the genotypes in Evita. Furthermore, it is straightforward to construct a neutral model for the E. coli evolution by building a system that copies the key E. coli parameters (population size, replication rate, mutation rate, genome size) and driving it with randomly selected births and deaths. So all that stands in the way of measuring evolutionary activity in such natural populations is the practical problem of gathering the raw data.

Evita and *E. coli* are both relatively simple evolutionary systems. Both, for example, are haploid. So one might worry that the evolutionary activity test cannot be applied in more complex contexts. But two considerations defuse this worry. First, those evolutionary systems are not as simple as they might seem at first. Evita, for example, has an indefinitely large space of genetic possibilities, since there is no limit to the length of Evita assembly language programs. And the assembly language is computationally universal, as well, so every algorithm whatsoever can be computed by some Evita string. But more importantly, the evolutionary activity test can be naturally extended to apply to many more complicated evolutionary systems. For example, if genetic novelty can be caused by recombination or other more complex genetic operators, this can be handled by including such operators in the neutral model. Or if the environment includes seasonal variation in resources or rare catastrophic changes, these can be added to the neutral model. Or if adaptation is reflected in something other than differential reproductive success, then it may well be able to be operationalized in some other way. How exactly to do so would depend on the details of the case at hand. It is impossible to explain how to handle all possible complications in advance, and there is no guarantee that the evolutionary activity test can be extended into all possible contexts. Practical problems might prove insurmountable. My point is that the test easily applies to simple cases and can easily be extended to many more complex ones.

A new defense of adaptationism

Adaptive explanations are the bread and butter of evolutionary biology. But the scientific legitimacy of such adaptive explanations is controversial, largely because of the classic paper “The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme” published by Stephen Jay Gould and Richard Lewontin in 1979. The controversy persists to this day in large part, I believe, because the fundamental challenge raised by Gould and Lewontin has not yet been met; in fact, it is rarely even acknowledged. The objectification of value in teleology described above can change this status quo, for we can now to defend the scientific legitimacy of adaptive explanations in a new and deeper way.

First, some terminology. I will refer to claims to the effect that a trait is an adaptation as an adaptive hypothesis. A specific adaptive hypothesis is a claim to the effect that a trait is an adaptation for some specified adaptive function, and a general adaptive hypothesis claims that a trait is an adaptation but identifies no adaptive function. A general adaptive hypothesis expresses the presupposition that the trait has some adaptive explanation. An example is the claim that large primate testes are an adaptation. An example of a specific adaptive hypothesis is the claim that large primate testes are an adaptation for producing more sperm. An adaptive explanation of a trait explains its existence or persistence as a result of adaptive evolution, i.e., by means of natural selection for that trait. Finally, by adaptationism I mean the thesis that the

activity of pursuing adaptive explanations of the existence and nature of biological traits is a normal and legitimate part of empirical science.¹²

Gould and Lewontin central complaint about adaptive explanations is that we have no principled way to tell when they are needed. People deploy adaptive explanations without justifying them over non-adaptive alternatives, such as appeals to architectural constraints or genetic drift. If one adaptive explanation fails it is simply replaced by another, but sufficient ingenuity enables any trait to be given an adaptive explanation. The general adaptive hypothesis that a trait is an adaptation is treated as untestable. As Lewontin puts it, “the adaptationist program makes of adaptation a metaphysical postulate that ... cannot be refuted” because the presupposition that a trait is an adaptation is never questioned (Lewontin 1977/1985, p. 76). The deeper worry is that the presupposition that a trait is an adaptation is really is untestable.¹³ There is a thicket of alternatives to adaptive explanations. How in principle can we tell when it is appropriate to pursue adaptationist branches? Gould and Lewontin summarize the predicament thus:

We would not object so strenuously to the adaptationist programme if its invocation, in any particular case, could lead in principle to its refutation for want of evidence. We might still view it as restrictive and object to its status as an argument of first choice. But if it could be dismissed after failing some explicit test, then alternatives would get their chance. (1979, pp. 258f).

The fundamental challenge, then, is to find some empirical test for general adaptive hypotheses. Without such a test, how can the practice of giving adaptive explanations be a normal and legitimate part of empirical science? In other words, the thesis of adaptationism would seem to be false.

Gould and Lewontin’s challenge to adaptationism provoked a storm of response. So many of the responses share the same basic form that this form can be called the “canonical” response. In a nutshell, the canonical response is to concede that there is no general empirical test for general adaptive hypotheses but construe this as on a par with normal empirical science.

Richard Dawkins nicely illustrates the canonical response when he considers traits that might not be adaptations. He points out that it is possible to test rival adaptive hypotheses by ordinary scientific methods, noting that “hypotheses about adaptation have shown themselves in practice, over and over again, to be easily testable, by ordinary, mundane methods of science” (Dawkins, 1983b, pp. 360f). Dawkins’s central point is that specific adaptive hypotheses have observable consequences, so they entail empirical predictions and thus can be tested. Dawkins illustrates this point with primate testes size. As it happens, primate testes size scales roughly but not exactly with body size. If testes weight is plotted against body weight, there is considerable scatter around the average line.

A specific adaptive hypothesis is that in those species in which females mate with more than one male, the males need bigger testes than in those species in which mating is monogamous or polygamous: A male whose sperms may be directly competing

with the sperms of another male in the body of a female needs lots of sperms to succeed in the competition, and hence big testes. Sure enough, if the points on the testis-weight/body-weight scattergram are examined, it turns out that those above the average line are nearly all from species in which females mate with more than one male; those below the line are all from monogamous or polygynous species. The prediction from the adaptive hypothesis could easily have been falsified. In fact it was borne out... (1983b, p. 361)

This illustrates how specific adaptive hypotheses can be tested by ordinary empirical methods.

Note, though, that Dawkins does not address the testability of general adaptive hypotheses. Furthermore, the test for specific adaptive hypotheses cannot be used to produce a test general adaptive hypotheses. The observable consequences of a specific adaptive hypothesis depend on the specific function hypothesized. Different functions may well entail different predictions. For example, the hypothesis that large primate testes are an adaptation for temperature regulation would entail a quite different prediction about where species fall in the testis-weight/body-weight scattergram. By contrast, the general hypothesis that large testes are an adaptation for something or other entails no prediction about where species fall in the scattergram. So, a general adaptive hypothesis inherits no observational consequences from specific hypotheses. For this reason, Dawkins admits that general adaptive hypotheses are untestable. "It is true that the one hypothesis that we shall never test is the hypothesis of no adaptive function at all, but only because that is the one hypothesis in this whole area that really is untestable" (1983b, p. 361). In other words, Dawkins thinks the fundamental challenge to adaptationism cannot be met. Of course, evidence for a specific function is a fortiori evidence for some function, so corroborating a specific adaptive hypothesis also corroborates the corresponding general hypothesis. But we cannot test all possible specific adaptive hypothesis for a trait. So the testability of specific hypotheses provides no test for general hypothesis.

The canonical response is weak.¹⁴ It capitulates in the face of Gould's and Lewontin's fundamental challenge by agreeing that there is no test for general adaptive hypotheses. But plenty of traits are not adaptations, and adaptive explanations are often inappropriate. Is there really no empirical way to tell whether adaptive explanations are in the offing? I think the answer is "Yes" and the objectification of biological teleology provides the key. The test we need is simply to collect and analyze evolutionary activity information.

The sign that an evolutionary process is creating adaptations is that its activity data are significantly higher than what would be expected if selection were random. If activity waves rise above the noise generated in a no-adaptation neutral model, then you know the corresponding items are adaptations even if you are ignorant about the adaptive functions. The activity data show that some adaptive explanation is needed even if it silent about the

merits of any specific explanation.¹⁵ In other words, the evolutionary activity method tests general rather than specific adaptive hypotheses.

Thus, the evolutionary activity test directly responds to the fundamental challenge to adaptationism. It parts company with the canonical response by not capitulating to Gould and Lewontin. As far as I know, it is the first response that takes this bull by the horns. The test does not assume that traits are adaptations but tests whether they are. Adaptive “just-so” stories have no place here; such stories propose specific adaptive hypotheses and these are not at issue. The issue is general adaptive hypotheses, and these are accepted only skeptically. Where the canonical response is weak, the activity test is strong. It makes the question of adaptation objective and empirical. When the adaptive stance is adopted, it is on the basis of empirical evidence against nonadaptive alternatives. So we can pursue the adaptationist program constructively and self-critically, as a normal and legitimate part of empirical science. Gould and Lewontin said that they “would not object to strenuously to the adaptationist programme if its invocation, in any particular case, could lead in principle to its refutation for want of evidence” (1979, pp. 258f). The evolutionary activity test provides just the sort of tool that Gould and Lewontin sought. So if we can take them at their word, they should now withdraw their objection.

Adaptationism is very controversial, so it is important to defuse some possible misunderstandings about the activity test. For example, some people equate adaptationism with the view that organisms are molded by their environment. However, the activity test does not presuppose this. The test applies perfectly well if the environment is significantly constructed by the organisms in the population. Furthermore, many who defend adaptationism also defend what could be called “panadaptationism” – the view that most (nonmolecular) traits of most organisms are due mostly to natural selection. The activity test takes no stand on panadaptationism. In particular, passing the activity test shows only that natural selection plays some role in the explanation of the trait in question. Other evolutionary forces like drift or architectural necessity might also be important factors in bringing about the trait; the point is that those non-adaptive forces are not sufficient to explain the trait if the evolutionary activity test is positive. Panadaptationists view evolution primarily as a process of adaptive hill climbing. Evolutionary activity measurements could actually be used to argue against this perspective. In effect, I sketch one such argument in the next section, where I suggest that the evolutionary activity signature of the continual production of evolutionary novelties cannot be explained merely by natural selection. That kind of evolutionary hyper-creativity depends on some more complicated mechanism than natural selection. We can start to appreciate this by comparing the evolutionary activity in biological and cultural evolution.

Towards a comparison on biological and cultural evolution

The information in evolutionary activity graphs can be summarized with statistics that reflect how evolution is creating adaptations. Such statistics have various uses, such as enabling quantitative comparisons of adaptive evolution across different systems. After informally explaining some of these statistics and explaining how they have been used to classify evolving systems, I will show how they shed new light on the relationship between biological and cultural evolution.¹⁶

When attempting to measure the degree of adaptive evolution in a system, one might try to reflect at least three different things. First, one might try to capture how well adapted the adaptations are, that is, how optimally they perform their function. Second, one might try to reflect the intensity of adaptive evolution, that is, the rate at which new adaptations are being produced by natural selection. Third, one might try to reflect the extent of adaptive evolution, that is, the total continual adaptive success of all the adaptations in the system. I will concentrate on the second and third ideas.

The intensity of evolutionary activity intuitively corresponds to the rate at new evolutionary activity is being created in the system, measured as the rate at which new activity waves are entering the activity graph. When there are very few new waves the intensity of activity is low; when a lot of new waves are being generated the intensity is high. To clean up this measure of the intensity of evolutionary activity, one would normalize the intensity observed in the target system with the intensity observed in a neutral model, yielding the excess intensity. One simple way to accomplish this is to measure not new waves but new waves that accrue more activity than would be expected in a neutral model.

The extent of evolutionary activity intuitively corresponds to the amount of evolutionary activity present in the system, measured as the sum total of the activity in an activity graph at a given time. If you think of the activity waves as being made up of grains of sand, the extent of activity at a given time is the mass of sand at that time in the graph, where the mass is weighted by its height in the graph. When the system has lots of very large activity waves, the extent of activity is very high. When the system has only a few waves and they are relatively low, the extent of activity is relatively low. As with intensity, one would clean up this measurement by normalizing the extent of activity observed in the target system with the extent observed in a neutral model, perhaps simply by subtracting the neutral extent from the target extent, thus yielding the excess extent.

The extent and intensity of evolutionary activity are two independently varying aspects of a system's adaptive evolution. For example, if adaptations continue to persist indefinitely without changing and no new adaptive innovations invade the system, then the extent of activity will continually increase, but the intensity of activity will fall to nil. On the other hand, if evolution is continually creating new adaptations and destroying older ones,

the intensity of activity will be positive but the extent of activity will be very low.

The intensity and extent of activity statistics are quite general and apply to data generated by both artificial and natural systems, and they apply at different levels of analysis. I have used evolutionary activity statistics to measure the creation of adaptations in a variety of evolutionary system (Bedau and Packard 1992; Bedau 1995; Bedau 1996; Bedau and Brown 1997; Bedau, Joshi, and Lillie 1999; Bedau, Snyder, Brown, and Packard 1997; Bedau, Snyder, and Packard 1998; Rechtsteiner and Bedau 1999a,b). Comparing data from a variety of different systems suggests that these statistics can be used to partition evolutionary dynamics into four qualitatively different classes. Class 1 consists of systems in which evolution creates no adaptations at all (e.g., all neutral models, systems in which the mutation rate is too high, and systems in which the selection pressure is too low, etc.). The signature for this class is zero excess intensity and extent of activity. Systems in which evolution has created adaptations but in which no new adaptations are being created fall into class 2 (e.g., stable ecosystems), with the signature of zero excess intensity and unbounded excess extent. Class 3 consists of systems that continually create new adaptations but are bounded in the amount of adaptive structure they contain (e.g., if new adaptations always supplant old adaptations). Its signature is positive excess intensity and bounded excess extent. If new adaptations are continually created and the total amount of adaptive structure continues to grow, then the system falls into class 4, which has the signature of positive excess intensity and unbounded excess extent. The biosphere as reflected in the fossil record exhibits class 4 dynamics. (For more details about this classification, see Bedau, Snyder, and Packard 1998, and Skusa and Bedau 2002.)

Class 4 is an especially explosive kind of evolutionary creativity. It is intriguing in part because no known existing artificial evolving system generates the same kind of behavior (Bedau, Snyder, Brown, and Packard 1997; Bedau, Snyder, and Packard 1998). Although we do not know the mechanism behind class 4 behavior, it seems to involve the course of evolution continually creating new kinds of environments that open the door to qualitatively new kinds of adaptations.¹⁷ There is some reason to think that a similar hyper-creative process might be at work in cultural evolution. We could start to assess this conjecture if we could apply evolutionary activity statistics to cultural evolution. I have recently started to do this in collaboration with Andre Skusa. Specifically, we have examined the evolution of technology as reflected in patent records, and we use evolutionary activity to create an empirical picture of the adaptive dynamics in patented inventions. Such pictures allow us to compare the dynamics of patented technology with those exhibited in biological evolution.¹⁸

Patents offer some important advantages for those looking for cultural evolution in empirical data. It is often difficult to operationalize the units of cultural evolution. It is difficult to distinguish new innovations from copies of old innovations when the items are ideas or other mental aspects of culture. Another difficulty is ascertaining precise genealogical relationships. One can

finer these difficulties by studying the evolution of technology as reflected in patent records. Although the evolution of inventions involves the diffusion and selection of ideas, one can identify individual inventions with individual patents. To be patentable an invention must meet three criteria: novelty, usefulness, and non-obviousness. So patented inventions are certified to be new and functional. A patent's novelty is documented by citing the previous patents (and sometimes published papers) that involve related ideas; these are called the patent's "prior art." The citations should identify all the important prior art from which the invention is derived, and in the aggregate they allow a patent's entire genealogy to be inferred.

The analogies and disanalogies between biological and cultural evolution are a matter of some controversy (Hull 1988, 2001) but it is relatively straightforward to extract evolutionary activity data from patent records. The units of evolution with which we are concerned (at least in the first instance) are individual patents; these are analogous to genes (or, as memeticists might suggest, "memes"). A gene could vanish forever from an evolutionary system. By contrast, a patented invention never goes fully extinct because the invention exists forever in the patent records. We consider that a patent "reproduces" when it leads to the production of other patents; that is, in contrast with most biological evolution, patent reproduction necessarily involves evolutionary innovation.

Especially successful or valuable patents tend to be those that are especially heavily cited. A large body of work in scientometrics has repeatedly confirmed that number of citations is a good reflection of the technological significance and economic value of a patented invention (Albert et al. 1991; Narin 1994; Pavitt 1985; Perko and Narin 1997; Albert 1998). Once a patent has received more than ten citations, the economic value reflected by each additional citation has been estimated to be more than one million US dollars (Harhoff et al. 1999). For these reasons our bookkeeping of an individual patent's evolutionary activity is based on summing the citations the patent has received. From this perspective, the adaptive success of a patented innovation is measured by the extent to which it spawns subsequent patented innovations. More specifically, we increment a patent's activity at a given time by the number of citations it receives from patents issued at that time.

The fact that a patent has received a few citations does not prove that the invention significantly shapes the evolution of subsequent inventions. A patent might be cited by one or two subsequent patents even if patents to cite were chosen entirely at random. As with evolutionary activity measurements in other contexts, we can evaluate a patent's adaptive significance by comparing its activity with the activity observed in a neutral model of patent evolution. Our patent neutral model mirrors a few key aspects of the real patent data. In both the same number of patents are issued each week, and they exhibit the same distribution into the various patent classes. Patent citations refer to the same number of pre-9/96 and post-9/96 patents, and the references to post-9/96 patents fall into the various patent classes according to the same distribution.

The key distinguishing feature of the neutral model is that the patents to be cited are always chosen randomly.

Figure 4 about here

Skusa and I studied the evolution of technology as reflected in the 868,535 utility patents granted by the United States Patent and Trademark Office between 9/96 and 7/02. Figure 4 shows the dramatic difference between the activity accrued by the most heavily cited real patents and the most heavily cited shadow patents. Overall, the citation levels of shadow patents are very much lower than the citation levels of patents with excess activity. This shows that high citation levels are not due to chance but reflect an invention's value. That is, an invention's salient "reproductive" activity is caused by selection for the invention because of its technological value.

Note that the activity of one patent in Figure 4 stands far above the rest, accruing almost twice as many citations as any other patent. This patent covers the technology that allows web browsers to display information such as advertisements while a page is being loaded a link is clicked. The second most heavily cited patent covers the technology that allows cell phones to receive email and faxes, and the third most heavily cited patent allows remote control of the receipt and delivery of wireless and wireline voice and text messages. All of the ten most heavily cited patents fall into the information technology sector, and seven of them involve the Internet.

More detailed information can be extracted from the evolutionary activity data (see Skusa and Bedau 2002). The point here is simply that evolutionary activity statistics make it feasible to visualize and quantitatively assess the adaptive evolutionary dynamics exhibited in cultural evolution. We have applied the method to technological evolution as reflected in patent record data, but it can be applied to a variety of other cultural systems. Our pilot project underscores the vast importance of information technology, and especially the Internet, over the past five years. This is not news, of course; it just corroborates what we already knew. But it does confirm the aptness and probity of evolutionary activity analysis of cultural evolution. Furthermore, it opens the door to quantitative comparison of cultural and biological evolution. And this provides a constructive empirical route for investigating whether the hyper-creativity exhibited by biological evolution also characterizes cultural change.

Conclusions

David Hull probably speaks for most philosophers and scientists when he says he wants to "avoid the use of such problematic notions as 'benefit'" in his treatment of natural selection because "their elimination from explanations of biological adaptations was one of Darwin's major achievements" (Hull 2001, p. 57). I want to counter this attitude with a gestalt switch. In my view, Darwin

does not remove value notions like benefit from our understanding of biological adaptation; he simply spells out the objective signs that reflect when those benefits play a certain kind of explanatory role. The gestalt switch I recommend is to treat Darwin's achievement not as the elimination of value in biology but as its objectification or operationalization.

I have tried to spell out some of the fruits of one case of this sort of objectification of value in biology. That we can objectify values in science is no surprise. We do it all the time. One objectifies the aggregate economic value of a country's commerce by measuring its gross national product, and one objectifies the social values of an individual or social group through appropriate public opinion surveys. Whether such objectifications of values in science are interesting depends mainly on the soundness of the objectification methodology and on the specific insights revealed (if any). The proof of this pudding is in the eating. If my case of objectification has any significant interest, it lies in the evolutionary activity statistics and their application in the new defense of adaptationism and the new ability to compare biological and cultural evolution. The main moral I draw from this one case study is that my value-centered view of teleology would never have born these fruits but for its objectification.¹⁹

References

- Adami, C., and Brown, C. T. (1994), Evolutionary Learning in the 2D Artificial Life System "Avida". In Brooks, R. and Maes, P., eds., *Artificial Life IV* (Cambridge: MIT Press), 377-381.
- Albert, M. B. 1998. *The new innovators: Global patenting trends in five sectors*. Washington D.C.: U.S. Department of Commerce.
- Albert, M., Avery, D., Narin, F., and McAllister, P. 1991. Direct validation of citation counts as indicators of industrially important patents. *Research Policy* 20: 251-259.
- Aunger, R., ed. 2000. *Darwinizing culture: The status of memetics as a science*. New York: Oxford University Press.
- Barkow, J. H., Cosmides, L., and Tooby, J., eds. 1992. *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.
- Beatty, John. 1987. Natural selection and the null hypothesis. In Dupré, J., ed., *The Latest on the Best, Essays on Evolution and Optimization* (Cambridge: MIT Press), pp. 53-75.
- Bedau, M. A. 1990. Against Mentalism in Teleology. *American Philosophical Quarterly* 27: 61-70.
- Bedau, M. A. 1991. Can Biological Teleology be Naturalized? *The Journal of Philosophy* 88: 647-655.
- Bedau, M. A. 1992a. Where's the Good in Teleology? *Philosophy and Phenomenological Research* 52: 781-806.
- Bedau, M. A. 1992b. Goal-Directed Systems and the Good. *The Monist* 75: 34-49.
- Bedau, M. A. 1993. Naturalism and Teleology. In S. Wagner and R. Warner, eds., *Naturalism: A Critical Appraisal*, Notre Dame, IN: University of Notre Dame Press, pp. 23-51.
- Bedau, M. A. 1995. Three Illustrations of Artificial Life's Working Hypothesis. In Banzhaf, W., and Eeckman, F. (eds.) 1995. *Evolution and Biocomputation – Computational Models of Evolution* (Berlin: Springer), 53-68. Available for download at <http://www.reed.edu/~mab/biocomputation.pdf>
- Bedau, M. A. 1996. The Nature of Life. In Boden, M., ed., *The Philosophy of Artificial Life* (New York: Oxford University Press), pp. 332-357. Available for download at <http://www.reed.edu/~mab/life.oxford.html>
- Bedau, M. A. 1998. Philosophical Content and Method of Artificial Life. In T. W. Bynam and J. H. Moor, eds., *The Digital Phoenix: How Computers are Changing Philosophy*, pp. 135-152. Oxford: Basil Blackwell.
- Bedau, M.A. 2002. A new defense of adaptationism. Paper to be presented at the Northwest Philosophy Conference, Lewis and Clark College, October 2002.
- Bedau, M. A., and Brown, C. T. 1997. Visualizing Evolutionary Activity of Genotypes. *Artificial Life* 5: 17-35. Available for download at http://www.reed.edu/~mab/vis_gtypes_alifejournal.pdf

- Bedau, M. A., S. Joshi, and B. Lillie. 1999. Visualizing Waves of Evolutionary Activity of Alleles. In A. Wu, ed., *Proceedings of 1999 Genetic and Evolutionary Computation Conference Workshop Program* (Orlando: GECCO Proceedings), pp. 96-98. Available for download at http://www.reed.edu/~mab/vis_gecco99.pdf
- Bedau, M. A., and Packard, N. 1992. Measurement of Evolutionary Activity, Teleology, and Life. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., eds., *Artificial Life II* (Redwood City, CA: Addison Wesley), pp. 431-461. Available for download at <http://www.reed.edu/~mab/alife2.pdf>
- Bedau, M. A., Snyder, E., Brown, C. T., and Packard, N. H. 1997. A Comparison of Evolutionary Activity in Artificial Systems and in the Biosphere. In Husbands, P. and Harvey, I., eds., *Proceedings of the Fourth European Conference on Artificial Life, ECAL97* (Cambridge: MIT Press), pp. 125-134. Available for download at <http://www.reed.edu/~mab/ecal97.pdf>
- Bedau, M. A., E. Snyder, N. H. Packard. 1998. A Classification of Long-Term Evolutionary Dynamics. In Adami, C., Belew, R., Kitano, H., and Taylor, C., eds., *Artificial Life VI* (Cambridge: MIT Press), pp. 228-237. Available for download at <http://www.reed.edu/~mab/alife6.pdf>
- Blackmore, S. 1999. *The meme machine*. New York: Oxford University Press.
- Brandon, R. N. 1990. *Adaptation and Environment*. Princeton: Princeton University Press.
- Boyd, R, and Richerson, P. J. 1985. *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Burien, Richard M. 1992. Adaptation: Historical Perspectives. In E. F. Keller and E. A. Lloyd, eds., *Keywords in Evolutionary Biology* (Cambridge: Harvard University Press) , pp. 7-12.
- Cavalli-Sforza, L. L., and Feldman, M. W. 1981. *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.
- Cooper, V. S., and Lenski, R. E. 2000. The population genetics of ecological specialization in evolving *Escherichia coli* populations. *Nature* 407, 736-739.
- Dawkins, R. D. 1982. *The extended phenotype: The long reach of the gene*. New York: Oxford University Press.
- Dawkins, R. D. 1983. Adaptationism was always predictive and needed no defense. *The Behavioral and Brain Sciences* 6: 360-361. Commentary on Dennett (1983).
- Dennett, D. C. 1983. Intentional systems in cognitive ethology: The "Panglossian paradigm" defended. *The Behavioral and Brain Sciences* 6: 343-355.
- Dennett, Daniel. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Simon and Schuster: New York.
- Dupré, J., ed. 1987. *The Latest on the Best, Essays on Evolution and Optimization*. Cambridge: MIT Press.

- Farmer, D, and Belin, A. 1992. Artificial Life: The Coming Evolution. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., eds., *Artificial Life II* (Redwood City, CA: Addison Wesley), pp. 815-840.
- Garfield, E., and Welljams-Dorof, A. 1992. Citation data: Their use as quantitative indicators for science and technology evaluation and policy making. *Science and Public Policy* 19: 321-327.
- Godfrey-Smith, P. 2001. Three kinds of adaptationism. In Orzack, S. H., and Sober, E., eds., *Adaptationism and Optimality* (Cambridge: Cambridge University Press), pp. 335-357.
- Gould, S. J., and Lewontin, R. C. 1979. The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society B* 205, 581-598. Reprinted in Sober, ed., *Conceptual issues in evolutionary biology* (Cambridge: MIT Press), pp. 73-90. References to this edition.
- Harhoff, D., Narin, F., Scherer, F. M., and Vopel, K. 1999. Citation frequency and the value of patented innovation. *Research Policy* 81: 511--515.
- Hull, D. L. 1988. *Science as a process: An evolutionary account of the social and conceptual development of science*. Chicago: University of Chicago Press.
- Hull, D. L. 2001. *Science and selection: Essays on biological evolution and the philosophy of science*. Cambridge: Cambridge University Press.
- Kimmura, M. 1983. *The neutral theory of molecular evolution*. Cambridge: Cambridge University Press.
- Lenski, R. E., and Travisano, M. 1994. Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. *Proceedings of the National Academy of Sciences, USA* 91, 6806-6814.
- Lewontin, Richard. 1977/1985. Adaptation. In R. Levins and R. Lewontin (1985), *The dialectical biologist* (Cambridge: Harvard University Press). Originally published as "Adattamento" in *Enciclopedia Einaudi*, vol. 1, edited by G. Einaudi (Turin, Italy, 1977).
- Lumsden, C. J., and Wilson, E. O. 1981. *Genes, mind, and culture*. Cambridge: Harvard University Press.
- Lynch, A. 1996. *Thought contagion: How belief spreads through society*. New York: Basic Books.
- Maynard Smith, J. 1978. Optimization theory in evolution. *Annual Review of Ecology and Systematics* 9: 31-56.
- Mayr, E. 1982. *The Growth of Biological Thought*. Cambridge: Harvard University Press.
- Mayr, E. 1983. How to carry out the adaptationist program? *American Naturalist* 121 (March 1983): 324-333. Reprinted in Mayr, E., (1988), *Towards a New Philosophy of Biology: Observations of an Evolutionist* (Cambridge: Harvard University Press), 148-159. References to this edition.
- Monod, J. 1971. *Chance and Necessity*. New York: Vintage.
- Narin, F. 1994. Patent bibliometrics. *Scientometrics* 30: 147-155.
- Orzack, S. H., and Sober, E., eds. 2001. *Adaptationism and optimality*. Cambridge: Cambridge University Press.

- Orzack, S. H., and Sober, E., 1994. Optimality models and the test of adaptationism. *The American Naturalist* 143: 361-380.
- Pavitt, K. 1985. Patent statistics as indicators of innovative activities: Possibilities and problems. *Scientometrics* 7: 77--99.
- Perko, J. S., and Narin, F. 1997. The transfer of public science to patented technology: A case study in agricultural science. *Journal of Technology Transfer* 22: 65--72.
- Raup, D. M. 1987. Neutral models in paleobiology. In Nitecki, M. H., and Hoffman, A., eds., *Neutral models in biology* (New York: Oxford University Press), pp. 121-132.
- Ray, T. S. 1992. An Approach to the Synthesis of Life. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., eds., *Artificial Life II* (Redwood City, CA: Addison Wesley), pp. 371-408.
- Rechtsteiner, A. and M. A. Bedau, 1999a, A Generic Model for Quantitative Comparison of Genotypic Evolutionary Activity. In D. Floreano, J.-D. Nicoud, F. Mondada, eds., *Advances in Artificial Life* (Heidelberg: Springer-Verlag), pp. 109-118. Available for download at <http://www.reed.edu/~mab/eca199.pdf>
- Rechtsteiner, A. and M. A. Bedau, 1999b. A Generic Model for Measuring Excess Evolutionary Activity. In Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., & Smith, R. E. (eds.), *GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference*, Vol. 2 (San Francisco, CA: Morgan Kaufmann), pp. 1366-1373. Available for download at <http://www.reed.edu/~mab/gecco99.pdf>
- Rogers, E. M. 1995. *Diffusion of innovations*. Fourth edition. New York: Free Press.
- Rosenberg, Alexander. 1985. *The Structure of Biological Science*. Cambridge: Cambridge University Press.
- Skusa, A. and M. A. Bedau. 2002. Towards a comparison of evolutionary creativity in biological and cultural evolution. In *Artificial Life VIII*, R. Standish, M. A. Bedau, and H. Abbass, eds., pp. 233-242. Cambridge: MIT Press.
- Sober, E. (1984), *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Cambridge: MIT Press.
- Sober, E. 1987. What is Adaptationism? In J. Dupré, ed., *The Latest on the Best, Essays on Evolution and Optimization* (Cambridge: MIT Press), pp. 105-118.
- Sober, E. 1993. *Philosophy of biology*. Boulder: Westview.
- West-Eberhard, Mary Jane. 1992. Adaptation: Current Usages. In E. F. Keller and E. A. Lloyd, eds., *Keywords in evolutionary biology*, (Cambridge: Harvard University Press), pp. 13-18.
- Wilson, E. O. 1978. *On human nature*. Cambridge: Harvard University Press.
- Wimsatt, W. C. 1987. False models as means to truer theories. In Nitecki, M. H., and Hoffman, A., eds., *Neutral models in biology* (New York: Oxford University Press), pp. 23-55.




Figure 1

Evolutionary activity waves in Evita (top) and its neutral shadow (bottom). Note that the activity scale on the neutral shadow is inflated by a factor of three, in order to highlight the neutral shadow waves (barely visible along the bottom).

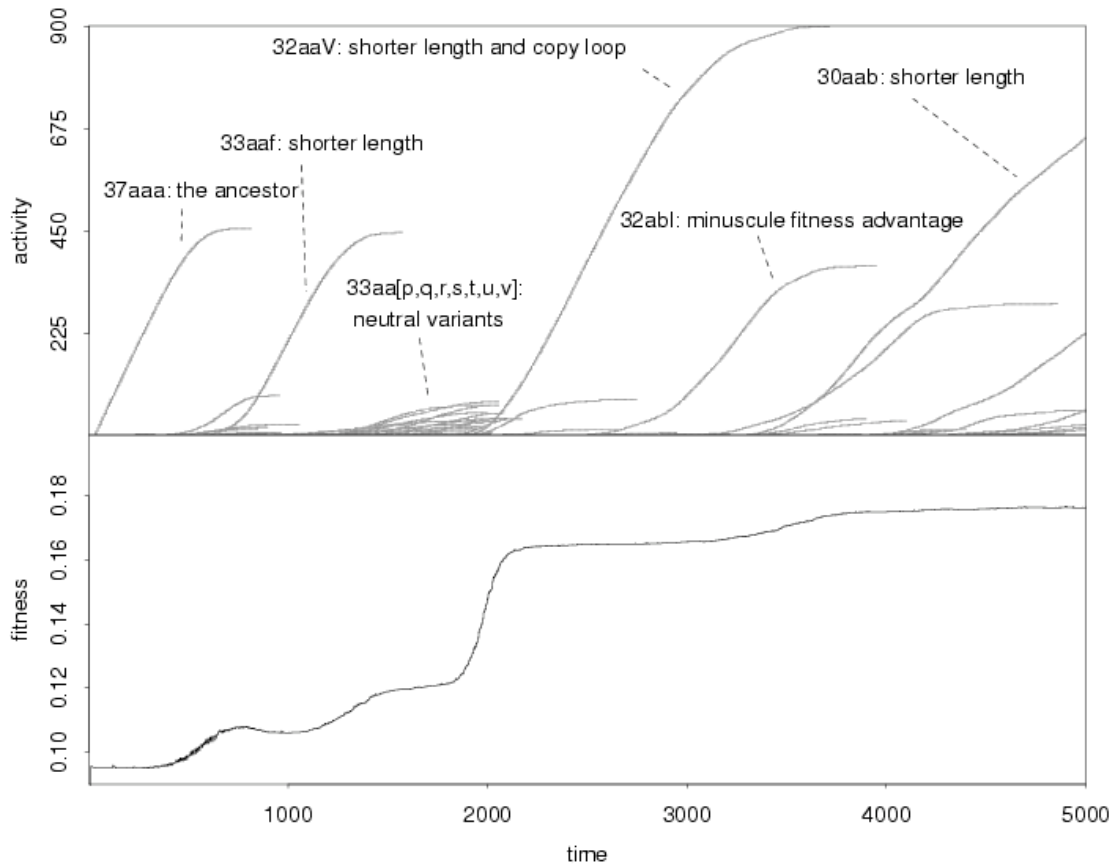


Figure 2

Evolutionary activity graph (above) and average population fitness (below) from a typical Evita run. The adaptive advantage of the genotypes causing the salient waves is indicated. The start of a significant wave generally corresponds to an increase in fitness. Note the cloud of neutral variants that cause one of the fitness jumps and act in the population like a single phenotype. These neutral variants are more fit than genotype 33aaf because they require one fewer instruction per execution of the copy loop. Note also that the significant wave due to genotype 32abl does not cause a significant fitness increase; it is nearly phenotypically equivalent to 32aaV because it executes only one fewer instruction per reproduction event than 32aaV.




Figure 3

Blow up of the evolutionary activity graph in Figure 2, showing the neutral variants that cause the fitness increase just after time step 1000.




Figure 4

Comparison of the activity of real patents and shadow patents, showing the twenty most heavily cited patents and the twenty most heavily cited shadow patents. The activity waves of the real patents all rise above 100 while those of the shadow patents remain below 25. Note that the activity accrued by significant real patents can vastly exceed that accrued by any shadow patent.

Notes

¹ Another interesting form of biological teleology has to do with the purported inherent “progress” exhibited by the course of evolution. For a discussion of how my objectification of value in biology bears on this, see Bedau 1998.

² For more on my value-centered view of teleology, and for comparisons with other views of teleology, see Bedau (1990, 1991, 1992a,b, 1993).

³ Not just selection of the trait; see Sober (1984) for this distinction.

⁴ “Selection test” information would be more informative (and more awkward) terminology. Norman Packard and I developed and applied this method to a number of systems over a number of years with the help of students and colleagues. See Bedau and Packard 1992, Bedau 1995, Bedau 1996, Bedau and Brown 1997, Bedau, Snyder, Brown, and Packard 1997, Bedau, Snyder, and Packard 1998, Bedau, Joshi, and Lilly 1999, Rechtsteiner and Bedau 1999a,b.

⁵ Although the evolutionary activity method is novel, the essential logic behind it should be familiar. See, for example, Kumura 1983, Wimsatt 1987, Raup 1987, Beatty 1987. These parallels are traced in greater detail in Bedau 2002.

⁶ Created by C. Titus Brown, Evita is inspired by Tierra (Ray 1992) and its derivative Avida (Adami and Brown 1994), but it is much simpler because it disallows the kind of interactions that lead to parasitism and the other interesting evolutionary phenomena observed in Tierra. Its simplicity makes it an especially simple and clear illustration of how graphing evolutionary activity reveals a system’s evolutionary dynamics. A much more detailed presentation of the material in this section is available elsewhere (Bedau and Brown 1997).

⁷ If the mutation rate is too low, there is no significant genetic change in the population. If the mutation rate is too high, the population dies out almost immediately because no successfully reproducing creature can survive the bombardment of mutations long enough to reproduce.

⁸ Actually, it has an indefinite number of them, due to random sampling differences—a qualification I will usually ignore.

⁹ In the Evita activity graphs shown here, the Evita system parameters were all identical except for mutation rate and elapsed time. Each genotype in a given run is given a unique name of the form Nxxx, where N is a number indicating the genotype's length and xxx is a three-character string (in effect, a base 52 number) indicating the genotype's order of origination among genotypes of that length. For example, 32aac is the third length 32 genotype to arise in the course of a given run.

The grid size was 40 x 40, so when the grid filled up the population consisted of about 1600 self-reproducing programs. I have pruned out irrelevant data about transitory genotypes by graphing only those genotypes that had at least five instances in the population at some time. This removes some of the “little hairs” created by nonadaptive noise (see Figure 3).

¹⁰ Note that the activity scale (y-axis) in these two plots is roughly comparable, except that activity on the bottom is expanded by a factor of three to make the neutral model activity easier to see.

¹¹ Notice that the fourth salient wave (due to genotype 32abl) does not correspond to a significant fitness jump. This genotype is well adapted, but it is not significantly better adapted than its main predecessor: genotype 32aaV. The waves from 32aaV and 32abl coexist for so long because the two genotypes are nearly neutral variants. In fact, the fitness of the second wave (32abl) exceeds that of the first wave (32aaV) by about only 0.5%. The interactions among the three salient waves between updates 4000 and 5000 have a similar explanation. They are a significant improvement (5% fitness advantage) over the genotypes that they drive extinct, but they differ from one another by much less (less than 2%).

¹² My use of the term “adaptationism” captures what I believe is the central issue, but I should emphasize that the term is sometimes used in other ways. To get a sense of the

similarities and differences, see Maynard Smith 1978; Dawkins 1982; Dupré 1987; Sober 1987; Brandon 1990; Burien 1992; West-Eberhard 1992; Orzack and Sober 1994, 2001; Godfrey-Smith 2001.

¹³ It might be uncontroversial that some specific traits are adaptations, but those are the exception.

¹⁴ One can begin to appreciate just how canonical this response is by examining some other well-known responses to Gould and Lewontin (e.g., Mayr 1983, Dennett 1983, Rosenberg 1985, Sober 1987, Dennett 1995, Sober 1993). I spell out these parallels elsewhere (Bedau 2002).

¹⁵ For example, inspection of Figure 2 shows that the big wave in the middle produced by genotype 32aaV is an adaptation but it does not show what makes it better than its peers. We can usually discover a genotype's adaptive advantage by independent microanalysis, as we did for 32aaV.

Genetic hitchhikers and genetic drift introduce some complications in this analysis. I discuss these details in a forthcoming monograph.

¹⁶ Mathematical details about the statistics can be found elsewhere (Bedau and Packard 1992; Bedau 1995; Bedau, Snyder, Brown and Packard 1997; Bedau, Snyder and Packard 1998).

¹⁷ The phrase "evolutionary novelty" is sometimes used to refer to the production of qualitatively new kinds of adaptations. So the evolutionary hyper-creativity of class 4 dynamics involves the continual production of new kinds of evolutionary novelties. Developmental systems theorists sometimes call attention to individual examples of evolutionary novelties. My focus here is complementary but broader: the general pattern of the continual production of new kinds of evolutionary novelties.

¹⁸ The material in this section is explained in greater detail in Skusa and Bedau 2002. There is plenty of previous work on cultural evolution and on patents, but none quite like ours. For many years cultural change has been treated as a process of the diffusion of ideas (Rogers 1995), and the scientometrics community has been investigating scientific and technological change by analysis of bibliometric data and patent records for decades (Pavitt 1985; Garfield, and Welljams-Dorof 1992; Narin 1994; Albert 1998). But these approaches think of "evolution" simply as any change in time rather than just change resulting from differential imperfect replication and selection. Sociobiology (Wilson 1978, Lumsden and Wilson 1992) and its contemporary sibling evolutionary psychology (Barkow, Cosmides and Tooby 1992) explore one kind of connection between biological and cultural evolution, specifically, the extent to which certain psychological and cultural phenomena (e.g., homosexuality and altruism) can be explained by appeal to the operation of biological evolution itself. This reduction of social science to biology is contrasted with the approach to culture illustrated by memetics (Lynch 1996, Blackmore 1999, Aunger 2000), which considers the evolution of cultural phenomena in its own right, independent from and even competing with biological evolution. The two classic quantitative treatments of cultural evolution (Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985) tend toward different answers to the question whether cultural evolution is ultimately explainable in terms of biological evolution, with Cavalli-Sforza and Feldman leaning toward explanatory dependence and Boyd and Richerson leaning toward a limited autonomy for culture. My approach is neutral on this issue. I study cultural evolution as an evolutionary process in its own right, ignoring whether and how it might depend on biological evolution. My goal is to provide an empirical and quantitative picture of the evolution of culture, one which allows us to compare its evolutionary dynamics with those of biological evolution. Both reductionists and antireductionists could profit from objective empirical measurement of cultural dynamics.

¹⁹ Special thanks to my collaborators on the evolutionary activity investigations reported here: Titus Brown, Norman Packard, and Andre Skusa. Special thanks also to James Lennox, for his helpful prepared commentary on the version of this paper presented at the Sixth Pittsburgh-Konstanz Conference in the Philosophy of Science. Thanks for helpful discussion to the audience in that conference, especially David Griffiths, Peter Machamer, Sandra Mitchell, and Jerome

Shaffner. For helpful comments or discussion on earlier versions of parts of this paper, thanks to Phil Anderson, Peter Godfrey-Smith, Mike Raven, Tom Ryckman, and Chris Stephens. Thanks also for helpful discussion to audiences at the University of Oklahoma, the University of Washington, Washington University in St. Louis, the Center for Humanities at Oregon State University, the Center for Cognitive Studies at Tufts University, the Santa Fe Institute, the first Genetic and Evolutionary Computation Conference, the fourth European Conference on Artificial Life, the sixth Artificial Life Conference, the Lake Arrowhead Conference on Computational Social Science, and at the Fraunhofer Gesellschaft in Sankt Augustin, Germany, where some of the material in the paper was presented. And thanks to the University of Oklahoma, its Zoology Department, and Professor Tom Ray, for hospitality and support while some of this work was accomplished.