

Where Guesses Come From: Evolutionary Epistemology and the Anomaly of Guided Variation*

EDWARD STEIN

*Dept. of Linguistics and Philosophy
Massachusetts Institute of Technology
Cambridge, MA 02139, USA*

PETER LIPTON

*Dept. of Philosophy
Williams College
Williamstown, MA 01267, USA*

ABSTRACT: This paper considers a central objection to evolutionary epistemology. The objection is that biological and epistemic development are not analogous, since while biological variation is blind, epistemic variation is not. The generation of hypotheses, unlike the generation of genotypes, is not random. We argue that this objection is misguided and show how the central analogy of evolutionary epistemology can be preserved. The core of our reply is that much epistemic variation is indeed directed by heuristics, but these heuristics are analogous to biological preadaptations which account for the evolution of complex organs. We also argue that many of these heuristics or “epistemic preadaptations” are not innate but were themselves generated by a process of blind variation and selective retention.

KEY WORDS: Evolutionary epistemology, heuristics, preadaptation, variation.

I. INTRODUCTION

Evolutionary epistemology is an approach to the theory of knowledge that is based on an analogy between epistemic development and biological evolution. Beyond this, however, the phrase “evolutionary epistemology” has been used by many people to mean a variety of things. One common interpretation of the phrase is that an evolutionary epistemology provides a biological account of belief. This is not the type of evolutionary epistemology with which we will be concerned here. Instead, we will focus on the sort of evolutionary epistemology popularized by Donald Campbell, an evolutionary epistemology which has as its “central insight . . . that biological natural selection, or some other [similar] selection process involving the real world, edits and selects among variations, providing the fit between belief or knowledge and the real world” (Campbell 1977, p.

30; Campbell 1974a). For Campbell, evolutionary biology does not itself provide a comprehensive theory of knowledge, but it provides a *model* for such a theory.¹ He begins his approach to evolutionary epistemology by looking at three examples of the problem of fit: biological evolution, visual perception and the development of scientific knowledge. How can we explain the fit between organism and environment? How can we explain the fit between our visual perception of the world and the way the world is? How can we explain the fit between scientific theories and the world they describe? According to Campbell, all three problems have the same solution: trial and error, natural selection, or, more specifically, blind variation and selective retention. All three “marvelous” fits can be best explained by natural selection.

The three major components of the model of natural selection are variation, selection and retention. In the modern Darwinian theory of biological evolution, genetic mutations provide the variations, the environment provides the selection, and reproduction provides the retention. The variations are not pre-designed. Rather, fit is achieved only through the hindsight of the selection process. In biological evolution, this amounts to saying that the mechanism of variation is not influenced by the effects the variations would have. The likelihood of a mutation is not correlated with the benefits or liabilities that mutation would confer on the organism. Rather, those organisms with features which make them less fit for survival do not survive in competition with other organisms in the environment which have features that are more fit. Evolutionary epistemology attempts to apply this blind variation and selective retention model to the growth of scientific knowledge and to human thought processes in general. In this essay, we will talk primarily about scientific knowledge, but what we say will also apply to general knowledge. We focus on scientific knowledge because, in science, the separation between variation and selection is clearer, and we want to focus our discussion only on variation and not selection.

In this essay, we will discuss just one of the major challenges facing evolutionary epistemology — how to explain the apparent guidedness of the variations that yield hypotheses that are candidates for testing. Since evolutionary epistemology attempts to establish a relation between biological and epistemic evolution, an immediate difficulty is that biological evolution seems to involve *blind* variation while epistemic evolution seems to involve *guided* variation. Genetic change is not influenced by the needs of the organism, but the production of new hypotheses by a scientist seems obviously influenced by the problem she is trying to solve. If the evolutionary epistemologist is going to argue for a strong analogy between biological and epistemic evolution, he must provide either an account of how epistemic variations seem guided but are in fact blind, or an account of how biological variations seem blind but are in fact guided. We will do

both, primarily arguing that epistemic variation is blinder than it seems, but also arguing that biological variation is more guided than it seems. We will explore two accounts of epistemic variation: the appeal to hidden chaos, which says that variations seem guided only because their underlying randomness is suppressed, and the appeal to epistemic preadaptation, which says that the variations are guided, in that they are restricted by heuristics, but that these heuristics are themselves retained from a process that was mostly random. In other words, guided variation comes from retention — the retention of heuristics which are produced by previous variations. We will argue that not only do these two accounts work together to characterize the actual blindness and the seeming guidedness of epistemic variations, but, further, that they are analogous to biological mutation and biological preadaptation. Hidden chaos is the epistemic correlate of biological mutation and epistemic preadaptation is the epistemic correlate of biological preadaptation. Just as biological mutations and biological preadaptations combine to explain biological variations, hidden chaos and epistemic preadaptation combine to explain epistemic variations.

Our discussion will proceed in the following manner. First, we will discuss variation in biology. We will note that biological evolution involves both strictly random variation and variation guided by preadaptation.² We will then discuss the parallels in epistemic variation: hidden chaos and epistemic preadaptation. In our discussion of epistemic preadaptation, we will consider the problem of innate heuristics — if we have certain biologically innate restrictions on our epistemic variation, evolutionary epistemology may be false. We will argue that such innate heuristics do in fact exist, that they are consistent with evolutionary epistemology, and that they do not occur in large enough numbers to trivialize evolutionary epistemology. We will conclude with a discussion of some objections to this view of epistemic variation.

II. BIOLOGICAL VARIATION

The details of biological variation help to explain the seeming guidedness of epistemic variation. Before Darwin's theory of natural selection was proposed, many people supposed that the amazing fit between organism and environment was due to the work of an intelligent creator. The argument from design used complex natural structures as evidence for the existence of a supernatural creator of the universe. The argument usually took the form of an analogy. Suppose you were walking along the beach and came upon a watch; after studying the watch, you noted its precise design, a design which could have only come to be if this watch were designed to work in a certain way. You would reason that this watch was created by somebody who intended it to fulfill a certain function, that of

keeping time. The argument from design compares the watch to the human eye. The eye, like the watch, is a precise, complex instrument that performs a particular function with great accuracy and dependability. How could the eye be as it is unless it was designed by an intelligent creator? Darwin, of course, had an answer: he argued that the process of evolution through natural selection, a process with no intentional or intelligent creator behind it, could produce an eye. Blind variation and selective retention is a mechanism that can do God's work. The claim that the amazing fit between organism and environment is due to concealed random processes is central to biology and to the analogy that evolutionary epistemology exploits.

Daniel Dennett provides a useful thought experiment which shows how fit can result from unseen randomness:

This thought experiment is inspired by one of the most elegant and insidious of the classic con games; it probably already has a name, but not knowing it, I call it the touting pyramid. You obtain a mailing list of serious gamblers, divide it in half, and send one half the prediction that team A will win the championship next week, and the other half the prediction that team A will lose. A week later, half your mailing list has received a true prediction from you — free of charge. Discard the other half of the mailing list; divide the remainder in half again, and send them a second brace of complementary predictions; this cuts down your pool of suckers, but now they have two “proofs” of your clairvoyance. After a few more “successes,” you announce that the free trial period is over; for your next prediction they will have to pay. (Dennett 1984, p. 93 footnote)

The same sort of phenomenon occurs in biological evolution: since we only see those organisms which survive, we are like one of Dennett's “suckers”. Because a sucker does not know about all the wrong predictions which are sent out, he thinks the con man knows in advance which team will win. Similarly, we think that evolution is purposeful, but it is not — nature is only guessing. Only when he looks at the larger picture can a sucker see that he is being conned; only when we look at the larger picture can we see how random variation is producing organisms which are fit.

Just as the pyramid theory defuses the argument that the con-man must know who will win since he made a string of correct predictions, Darwin's theory defuses the argument from design. The theory of natural selection defuses the inference from design to designer not just because it provides an alternative explanation of design and so shows that the inference to a supernatural designer is not obligatory, but also because Darwin's theory is more plausible than the argument from design. Darwin's theory does not require positing things for which we have no good evidence. The argument from design involves a being of unimaginable powers, while the theory of natural selection involves forces that we observe today, forces that are easily explained.

In addition to the virtue of greater plausibility that comes from appealing to more mundane entities and processes, Darwin's theory is more plausible because it is more explanatory than the appeal to God. The appeal to God has an explanatory weakness shared by all straightforward intentional explanations: the pattern to be explained is used in the explanation as an intentional object. This does not make such explanations worthlessly circular, but it does make them fairly superficial. The design in nature is explained by appealing to that same design in God's mind. By contrast, as Robert Nozick has observed, explanations like Darwin's go deeper, because they can account for a pattern without appealing to that pattern itself in any form (Nozick 1974, p. 18–19).

The theory of natural selection does, however, face various anomalies. One of these, which will be of special importance to our discussion of evolutionary epistemology, is the problem of complex organs. It is virtually impossible for a single mutation to occur that would produce a complex organ such as a wing, since many different genes are involved. The probability that the necessary number of mutations would occur simultaneously, in a coordinated fashion, is infinitely small. So it appears that the natural selection model cannot account for the evolutionary development of these organs.

This anomaly can be resolved by what biologists call 'preadaptation.'³ In answer to the question "How do we get a complex structure like a wing?", biologists say that there are simpler intermediate structures which occurred before the wing. But this does not seem to solve the problem, unless these intermediate structures themselves have some adaptive advantage. If the intermediary structure had no adaptive function, then it would generally tend not to be retained and so could not lead to a wing. Obviously, the half-wing could not enable its possessor to fly (otherwise, it would be considered a wing). There must be an adaptive advantage of a pre-winged bird having a half-wing structure. The biologist's answer is to give an account of the half-wing's preadaptive function. For example, the half-wing may have been used for trapping insects. For biologists, the move to preadaptation allows them to explain how a complex organ may have evolved; it does so by claiming that the structure evolved from an "ancestor" of the organ which may have had a different function.

In summary, biological variation primarily involves two sorts of explanations for how features of organisms come into being: regular random mutations and preadaptations. Note that neither explanation alone will suffice. Random mutations without biological preadaptation do not provide a good explanation for complex organs since the chances of mutations alone producing complex organisms are infinitesimally small. Biological preadaptation without random mutations is no different from the argument from design — without random mutations behind them, the preadaptations themselves would be inexplicable.

III. EPISTEMIC VARIATION

Like its biological inspiration, evolutionary epistemology faces various anomalies. Perhaps the most conspicuous is that epistemic variation — the production of new hypotheses — seems anything but random. If scientific conjectures were really random, it would be a miracle that the astronomer who generates a new conjecture usually manages to produce a candidate hypothesis about asteroids rather than about steroids. And, of course, the scientist does much better than merely getting the field right: the variation in science appears tightly guided. This seems a dramatic disanalogy to the blindness of biological variation, thus calling evolutionary epistemology into question. Our view is that the claimed disanalogy does not exist. The appearance to the contrary rests in part on an overly simple picture of biological evolution, a picture which does not take into account the enormous role of biological preadaptation. We will argue that evolutionary epistemology can account for the apparent guidedness of variation, and that the central mechanism of that explanation is itself analogous to the mechanism of biological preadaptation that saved Darwin's theory from the anomaly of complex organs.

The remainder of this paper, which contains our explanation for the apparent guidedness of epistemic variation, falls into five parts. First, we consider the extent to which this guidedness is only apparent: the variation is in fact blind or chaotic, but, for various reasons, the chaos is hidden. Second, we show how the guidedness that actually exists can be accounted for by appeal to epistemic preadaptations in the form of heuristics and we discuss some of the forms these heuristics take. Third, we consider the mechanisms by which these heuristics themselves were generated. We will argue that most of them are based on hidden chaos, (i.e., that these heuristics have been retained from earlier blind variation) but some may be innate, and we will consider how an evolutionary epistemology ought to treat this innate component. Fourth, we will evaluate the strength of the analogy between biological and epistemic preadaptation by giving a positive account of it. Finally, we will consider some objections to the analogy between biological and epistemic variation.

Our first response to the seeming guidedness of epistemic evolution is to claim that it is, like the seeming guidedness of biological evolution, only apparent: both can be explained by unseen randomness of the sort exemplified in the touting pyramid. The analogy is as follows. Before Darwin, the common intuition was that the amazing fit between species and the environment is the result of a guided, intelligent process; after Darwin, species are seen to be the result of random variation coupled with selection and retention. Similarly, in the case of epistemic evolution, before evolutionary epistemology, the common intuition was that ideas and scientific discoveries are the result of guided intelligent processes;

after evolutionary epistemology, they are seen to be the result of random variation.⁴ The first response thus involves a straightforward appeal to hidden chaos — the variation is blind but the blindness is hidden. In the touting pyramid, the chaos is hidden because the suckers do not receive any of the wrong predictions — they do not yet see that the con man is not clairvoyant. In biological evolution, the chaos is hidden because we do not see all of the mutations (many of them fatal) that took place hundreds or thousands of years ago. In epistemic evolution, according to the appeal to hidden chaos, the chaos is hidden for a variety of reasons.⁵ Here we will only sketch some of the ways that it could possibly be hidden.

In order to see three of the ways that the randomness could be hidden, we will consider the hypothetical case of a scientist who has developed a theory T about the world. If we watch this scientist, we will think that she came to believe in T by an intelligently guided process. This, however, may only be an illusion. First, there may have been some unconscious variation going on in her mind. She may have unconsciously thought of many different theories, but only one enters her consciousness. Second, she might have consciously considered many alternative theories before she settled on T, but forgotten these pondered alternatives. The chaotic speculation would be hidden from her because she only remembers thinking about T. Third, it might be that she came upon T as a result of an accident or a lucky guess, but that afterwards she, for the sake of her peers, made up justifications for it, thereby hiding the fact that T was randomly derived. The random roots of T would thus be hidden behind her *post hoc* explanations. These are just three of the different ways that the randomness of variation might be hidden to the observer. Like Dennett's touting pyramid, they provide accounts of what is happening behind the scenes that makes a random process seem guided; as with Dennett's con man, the scientist's bad guesses may be suppressed.

But these simple appeals to hidden chaos are not enough. Earlier we discussed the anomaly of complex organs that occurs on the simple account of biological evolution — it is highly implausible that a single set of simultaneous mutations should produce such complex structures. A similar anomaly occurs in evolutionary epistemology. It is implausible that a single instance of hidden chaos could produce such a complex hypothesis as a detailed scientific theory. It seems wild to suppose that Kepler generated his laws through simple hidden chaos. Just as random mutations alone do not explain complex organs, hidden chaos alone does not explain our complex beliefs. Inspired by the biologist, the evolutionary epistemologist resolves this anomaly with an appeal to epistemic preadaptation. Epistemic preadaptations may act to guide the variation that yields hypotheses. Like the half-wing of biological preadaptation, the epistemic preadaptation had to be good for things other than its current adaptive use (or, at least, not harmful), otherwise, it would not have been adopted in

the first place. Just as the half-wing was a pre-wing structure which was good for something other than flight, so the epistemic preadaptation is good for something other than what the hypothesis it precedes is good for.

The epistemic variation that preadaptation is supposed to explain concerns the production of candidate hypotheses for testing. When these hypotheses are not generated randomly, we can call whatever guides or restricts their production a *heuristic*. Of course not all heuristics govern the production of hypotheses: we rely on them in most of our other epistemic activities as well. For example, we have heuristics that guide us in the testing of hypotheses and others that guide us in the acquisition of new hypotheses that require no testing since they have been deduced from previous hypotheses. But the heuristics that concern us in this essay are just those that guide us in acquiring hypotheses that we will then go on to test, since they are the ones whose existence seems to call the analogy to biological variation into question. In science, these are the heuristics that operate in the “context of discovery”, where conjectures are found, as opposed to the “context of justification”, where they are tested. Accordingly, we call these heuristics “discovery heuristics”. The two major questions about discovery heuristics concern their nature and their source. We will discuss these in turn.

As a hypothetical example, consider a chemist who is trying to discover why a particular compound behaves in a particular way. In her search, she would not make a random conjecture and start to refute it, as might be predicted by an evolutionary epistemology relying only on the appeal to hidden chaos. Rather, she would draw on a panoply of heuristics which she has developed in part from her previous chemical experience. Perhaps she would consider explanations for the behavior of other similar compounds, drawing her conjecture from these. In doing so, she is making certain assumptions: she is assuming that the other compound is in fact similar, she is assuming induction (namely, that past experiences with this and similar compounds will be repeated in the future), etc.

Given our broad notion of a heuristic as anything that restricts variation, there should be many specific types of discovery heuristics. In this essay we can only briefly discuss three: abstract rules, concrete rules, and exemplars. Abstract rules are discovery heuristics that apply to many areas of inquiry. Some of these rules are explicit, in the sense that they are ones a scientist could be expected to volunteer or at least to acknowledge as her own as soon as they are presented. The actual content of particular discovery heuristics is not important to our argument, but perhaps some examples are that hypotheses should be potential solutions to the problem at hand, that they should be tractable with the tools of the discipline, and that they should be testable. These explicit rules may seem vague or vapid, but they are clearly highly restrictive as compared to strictly random variation. There are presumably also many general rules with more

articulated and restricting content, but most of these must be implicit or tacit, since scientists do not volunteer them, psychologists have not discovered them, and philosophers cannot agree about them. For example, there are presumably various abstract features of a hypothesis that give it the sort of simplicity and coherence that help to make it a promising conjecture, and the scientist may be supposed to have tacit rules that guide him to produce conjectures with these features.

Unlike abstract rules, concrete rules are field-specific discovery heuristics. They tell scientists to favor conjectures that bear certain relationships to other theoretical claims in the discipline. Many concrete rules can be seen as rules for modifying old hypotheses to produce new ones. When a previously established hypothesis is judged to be unacceptable, it is rarely rejected without a trace. Typically, the old hypothesis provides a strong constraint on the form that its replacement will take. In the simplest case, the rule would be that the new hypothesis should be as similar to the old one as the data will allow. The same sort of rule often appears to be applied in larger scale theory change: the new theory should not only entail the old data, it should also entail an approximation of the old theory.⁶ To make the contrast with abstract rules more vivid, we may say that it is the old hypothesis itself that provides the concrete rule which acts to restrict variation, though the scientist can only be expected to have tacit knowledge of the many ways the old hypothesis constrains the new one.

The last type of discovery heuristic we will discuss has been emphasized by Thomas Kuhn (Kuhn 1970, *passim*, especially the postscript). A scientist uses a canonical solution in her field as a concrete model for new solutions. This exemplar focuses the scientist's attention on new putative solutions that she perceives as similar to the exemplar. The similarity relations the exemplar creates sharply restrict the variation the scientist undergoes in solving her problem. Indeed, if Kuhn is right, they also determine what problems she will attempt to solve.⁷ Unlike abstract rules, exemplars are highly field-specific; unlike concrete rules, exemplars are particular solutions that remain entirely acceptable rather than general hypotheses or theories that usually need to be modified. Like the two other types of discovery heuristics, exemplars partially account for the directed character of epistemic variation.

Having given a sketch of the nature of some discovery heuristics, we turn now to their source. Although these heuristics help to explain why epistemic variation is restricted, heuristics alone cannot provide a complete explanation. Just as biological mutations and biological preadaptations must work together to explain away the seeming guidedness of biological variation, hidden chaos and epistemic preadaptation must work together to explain away the seeming guidedness of epistemic variation. In biology, the half-wing appeared as a result of a random genetic mutation. Similarly, a non-innate preadapted heuristic must be the result of hidden

chaos, otherwise it seems that evolutionary epistemology will not work. Epistemic preadaptation thus involves admitting that not every conjecture is completely unguided; some are guided by preadapted heuristics, generalized rules of thumb developed from past experiences. The non-randomness of some acts and some scientific discoveries is attributed to heuristics which are themselves retained from previous blind variation and selective retention and can be considered “already achieved wisdom of some general sort” (Campbell 1974a, p. 422). Thus we must distinguish the question of the source of our hypotheses from the question of the source of the heuristics that help to generate them. Because of heuristics, the generation of a hypothesis may not be random or blind; but this leaves open the possibility that some of the heuristics come from blind variation.

As a hypothetical example, consider a scientist who makes a conjecture C guided by a certain heuristic H_1 . Where did this scientist come up with H_1 ? Evolutionary epistemology says that H_1 , if not innate, must be the result of a blind variation and selective retention process. However, there remain two possibilities: H_1 might have been actually blindly selected, or H_1 might have been selected on the basis of another heuristic H_2 (a meta-heuristic). The same is true with H_2 — there can be a nested hierarchy of heuristics — and there seems to be a possibility of an infinite regress, a regress which must come to an end if evolutionary epistemology is going to work. The regress can be prevented if the appeal to epistemic preadaptation is qualified; the qualification is that, ultimately, behind every instance of preadaptation not based on an innate heuristic, there is a case of hidden chaos. In other words, each non-innate conjecture, though it may be based on many levels of heuristics, must ultimately be based on a blind guess. Thus, we see that the two responses to the challenge against the blindness claim (hidden chaos and epistemic preadaptation) come together. A seemingly guided discovery is attributed to blind variation in one of two ways: either the discovery is viewed as a result of hidden chaos, or as a result of guided processes based on heuristics retained from previous variation, ultimately stemming from hidden chaos.

In the last two paragraphs, we have qualified evolutionary epistemology by restricting our discussion to non-innate heuristics; we leave open the possibility of there being innate heuristics. If, however, there are certain heuristics that are innate, then it looks as though evolutionary epistemology is false, since there would then be heuristics that do not have their source in blind variation. For example, if Noam Chomsky is right in claiming that we have an innate capacity to learn a certain type of language, then the acquisition of language would be based on a set of innate heuristics, principles genetically programmed into all humans. Similarly, it may be that some heuristics that govern scientific practice, such as some sort of inductive heuristic, are innate. Many evolutionary epistemologists have been tempted to respond to this problem with an appeal to blind variation

and selective retention at the *biological* level. They argue that epistemic variation can still be blind even when preadaptation traces back to an innate trait in humans if the innate trait can be attributed to the blind variation of biological evolution. As an example, let us consider an imaginary dialogue between an evolutionary epistemologist (EE) and one who is skeptical of evolutionary epistemology (SK).

EE: Knowledge processes and the growth of science are instances of blind variation and selective retention.

SK: But when I put words together to make sentences, there is no blind groping going on. In order to make myself understood, I form sentences without any blind variation.

EE: That is because you have developed, through a process that I call epistemic preadaptation, heuristics for communicating, for forming sentences. These preadapted heuristics were arrived at through a trial and error process. At first, you tried to communicate without a regular grammar, but you found that no one did what you wanted them to do. Through trial and error you found that if you asked for something in a certain way, you were more likely to get it. So, while you do not go through the trial and error process each time you speak, you do use techniques (heuristics) which you acquired through blind variation and selective retention.

SK: But what if the ability to use and acquire language is innate in humans? If this is the case, then there is no trial and error in the ability to acquire and use language. A set of innate principles of language is genetically programmed or “hard-wired” into humans.

EE: But this ability became hard-wired in humans through genetic evolution. The innate ability to acquire language, if there is such an ability in humans, came to be innate through genetic mutation and natural selection.

This line of thought seems to save evolutionary epistemology because biological evolution brings in the blind variation and selective retention that is needed to save the analogy. However, we think that this appeal to the blindness of biological variation creates a disanalogy between biological and epistemic evolution and makes evolutionary epistemology trivial. First, this appeal creates a disanalogy between epistemic and biological evolution. It pictures biological evolution as blind on its own level, while epistemic evolution is not blind in its own terms, but only blind by proxy, because it rests on the truly blind biology. By using biology both as an analogy and as a partial direct explanation, the strength of the analogy is weakened. Second, this appeal reduces the claim of evolutionary epistemology simply to the claim of the truth of Darwinian evolutionary theory. For if the appeal to biological variation is a legitimate one, then regardless of what occurs at the epistemic level, if biological evolution is

blind, then so is epistemic evolution. Consider, for example, the possible world in which biological evolution is true, but all our beliefs⁸ (that is, all the beliefs that we will ever have) are innate. In this possible world, all of our thoughts are guided (in fact, they are preprogrammed), while all biological evolution is blind, non-guided, etc. The appeal to the blind variation of biological evolution can still be made; the only difference would be that all the beliefs of these completely guided beings are like the innate capacity for language which we are claimed to have. In other words, if the appeal to biological variation is a legitimate way of saving the blind variation claim of evolutionary epistemology, then none of the other claims of the theory is needed; hidden chaos and epistemic preadaptation need not be true for evolutionary epistemology to be true. We do not, however, want to count an epistemology as evolutionary if its only source of blindness is biological variation. There are two versions of evolutionary epistemology involved here. The one we are defending claims that epistemic evolution is analogous to biological evolution. This is a strong and interesting claim, particularly if the analogy is spelled out, as we are here attempting to do, at least partially. The second version, the version on which the appeal to biological variation is based, claims that epistemic evolution comes *from* biological evolution. In the sense that this claim is normally understood, it is epistemically empty: it amounts only to the biological claim that we are a product of natural selection. Therefore, the appeal to biological variation is not a legitimate way to save evolutionary epistemology from attacks against the blindness claim because it trivializes evolutionary epistemology.

We are not disturbed by the trivializing result of the appeal to biological variation because we think that this move is not necessary to preserve the analogy between biological and epistemic evolution and to handle the objection that some heuristics may be innate. In fact, the distinction between innate and acquired beliefs is, we think, the saving grace of evolutionary epistemology. To begin with, evolutionary epistemology has no responsibility to explain genetic traits of organisms. Evolutionary epistemology just involves the claim that beliefs evolve in the same way that biological species do; innate heuristics need not be explained by evolutionary epistemology. It is perfectly consistent for an evolutionary epistemology to leave the explanation of such innate heuristics to biology; evolutionary epistemology only has to explain those heuristics which are *acquired*. Further, evolutionary epistemology need not deny that acquired beliefs are often partially rooted in innate heuristics — this does not mean that epistemic variation is not predominantly blind. The upshot of this is that heuristics can have a “mixed parentage”; they can be based on blind” epistemic variation as well as genetically innate heuristics.

What does this clarification of evolutionary epistemology say about epistemic preadaptation? Recall that the appeal to epistemic preadapta-

tion was made in order to save the claim that epistemic variation is blind. On the clarified view, this move is still necessary and legitimate. The appeal to biological variation was made in order to save the claim that epistemic variation is blind even in light of innate beliefs. On the clarified view, this move is not necessary, because it is no curse or surprise that some heuristics and hypotheses are linked to innate characteristics; the acquired component of epistemic variation can still be seen as blind.

Chomsky's account of language provides a good example of mixed parentage. The universal grammar provides the elements and principles that are common to all possible human languages. Chomsky (1986) distinguishes between an internalized language (I-language) and a target language. A target language is the language actually spoken while an I-language is some element in the mind of a user of a language. Chomsky has switched the search for the universal grammar from the target language to the I-language. He interprets universal grammar "as the theory of human I-languages, a system of conditions deriving from the human biological endowment that identifies the I-languages that are humanly accessible under normal conditions" (Chomsky 1986, p. 23). The universal grammar is "a characterization of these innate, biologically determined principles, which constitute one component of the human mind — the language faculty" (Chomsky 1986, p. 24). In our terms, the universal grammar, the innate capacity for language, is an example of a set of innate heuristics. The universal grammar restricts epistemic variation — certain "languages" are not humanly possible because of this restriction. The set of innate heuristics plus non-innate heuristics, acquired rules of language, together produce the actual language that we speak, our target language. The I-language isolates the innate component of language, but acquiring an target language requires setting particular parameters within the constraints of the I-language. Since the innate component does not *determine* the target language, it remains open to the evolutionary epistemologist to claim that the particular target language comes from selection from among those target languages which are compatible with the I-language. The example of universal grammar thus provides an excellent example of the innate heuristics and the "mixed parentage" of certain epistemic entities.

Our view, then, is that discovery heuristics may be either innate or randomly generated and, if innate, the evolutionary epistemologist (*qua* EET-type evolutionary epistemologist) has nothing to say about them. This avoids the trivialization of evolutionary epistemology, but it may create a new problem. We do not want to purchase the benefit of falsifiability at the cost of falsification. If all heuristics are innate, this would seem to refute our version of evolutionary epistemology, or at least make it vacuous. One reply would be to argue that even if all heuristics are innate, there is still a mixed parentage that allows for evolutionary

epistemology, since the heuristics only determine the range within which conjectures will fall and variation within that range is random.⁹ After all, different scientists make different conjectures. This reply would give evolutionary epistemology some role, but it would be disappointing as the sole response, since it would fail to meet the objection that epistemic variation is disanalogous to biological variation, the objection that motivates this essay. Of course, this is primarily an empirical question: what percentage of our heuristics are innate? If all our heuristics are innate, then evolutionary epistemology (at least the EET-type which we are doing) would be falsified.¹⁰ Fortunately, a stronger reply is possible: not all our heuristics are innate; many of them rest on hidden chaos. Below, we present some considerations that suggest that *most* of our heuristics are *not* innate.¹¹

One argument for acquired heuristics rests on underdetermination. Consider the set of all heuristics and beliefs that could plausibly be innate. For example, certain beliefs about how to produce a grammatical sentence are plausibly innate, but beliefs about phlogiston or black holes are not. Add to this set all the actual beliefs that could be plausibly held to be observational. This set of innate and observational beliefs underdetermines the conjectures scientists produce; in other words, the actual conjectures which scientists produce are not entailed by this set. Moreover, this set also underdetermines the heuristics, in our broad sense of the term, that scientists use. Scientific conjectures are much more severely constrained than they would be if the innate constraints provided the only restriction. It follows that many heuristics are not innate.

Another related reason why it seems so likely that there are acquired heuristics is that so many heuristics are field specific, like the concrete heuristics and the exemplars discussed above. No one supposes that particular hypotheses, theories, or exemplars are innate or entailed by the data, and these are also heuristics. Similarly, since these sorts of heuristics change over time within a single field, they are almost certainly not innate. Our beliefs and heuristics change much more rapidly than do our genes.

Thomas Kuhn's analysis of the role of exemplars in science makes a particularly vivid case for the claim that many heuristics rest on blind variation (Kuhn 1970, *passim*). According to him, new exemplars are generated when a scientific discipline goes into a protracted period of crisis. If the old exemplars cannot be made to work, some scientists will begin to search for replacements, a search which may culminate in a scientific revolution. But insofar as it is exemplars that guide scientists, and new exemplars can only be discovered by suspending the old ones, the search for new exemplars must be largely unguided.¹² Moreover, should the search be successful, this unmethodical mechanism will be hidden from view. After a revolution, much old research is re-presented as flowing from the new exemplars, a practice that gives the illusion of

cumulative scientific history based on fixed heuristics and so hides the chaotic source of the current exemplars. In short, new exemplars must rest on hidden chaos.

Having explained how the guidedness of epistemic variation is explained by hidden chaos and epistemic preadaptation as well as some innate heuristics, we will now step back and consider the strength of the analogy between biological and epistemic evolution. We have already extensively discussed the extent to which both involve the blind variation and selective retention mechanism — this is the most central feature of the analogy. We have also discussed the analogous anomalies that appear in biological and epistemic evolution. Biological variation cannot offer a good explanation of complex organs because the chances of strictly random mutations producing such apparent fit are so small. Similarly, epistemic variation cannot offer a good explanation of complex theories because the chances of strictly hidden chaos producing such apparent fit are also quite small. Further, we have seen that biological and epistemic evolution can solve their respective anomalies by appealing to preadaptation. In the biological case, a complex organ evolves from an earlier structure that was itself adaptive; in the epistemic case, a complex conjecture evolves from an earlier conjecture which we call a heuristic. In both cases, an infinite regress is avoided because the preadaptations are ultimately based on some truly random occurrence: in the biological case, a mutation was behind the preadapted structure and in the epistemic case, hidden chaos was behind the preadapted heuristic. Note that in both cases, the chaos is doubly hidden — it is hidden originally through the hidden chaos, such as that of the touting pyramid phenomenon, and then it is further hidden by the levels of preadaptation.

The similarity of biological preadaptation and epistemic preadaptation is underscored by our discussion of exemplars. Exemplars can be thought of as old solutions grafted onto new problems. In this way, they are much like biological preadaptations — old organs, which perform a function in the old environment, are used for another task in the new environment. In addition, exemplars determine what problems a scientist will tackle (namely those which are similar to the exemplar problems) in a way very much like the way a pre-organ structure restricts the sort of final organ an organism will have.

An additional feature of the analogy is that both types of preadaptations can be viewed in two, quite compatible, ways: either as restrictions on future variations or as programs for the generation of new variations. In both the biological and the epistemic cases, preadaptations restrict variation, but they also make certain variations possible. In biological evolution, the appearance of the half-wing makes possible the appearance of some structures, most notably a wing. At the same time, the appearance of the half-wing prevents certain other structures from appearing — for

example, an arm in place of the half-wing. In epistemic evolution, if a person adopts a certain heuristic, then various future conjectures are made possible, but others are made impossible or overwhelmingly unlikely.

One possible disanalogy between biological and epistemic variation that seems to have appeared can be turned into another analogous feature between the two. In biological evolution, all adaptations are ultimately based on strictly random mutations. In contrast, in epistemic evolution, some (though if our argument above is correct, relatively few) of the preadaptations are innate. The possibility of innate heuristics seems to create a disanalogy, since there does not on the surface seem to be a biological correlate to an innate epistemic preadaptation. This disanalogy dissolves if we recall that a heuristic is anything that guides the production of future theories. Carrying this view of preadaptation back to biology, we see that just as there may be certain innate heuristics that act to guide epistemic variation, there are also innate restrictions on biological variation, what we call genetic heuristics. The notion of genetic heuristics can be seen in two ways, depending on whether we look at phenotypic variation or genotypic variation. If we look at genotypic variation, we see that the structure of genes and DNA act as genetic heuristics. Certain biological variations are precluded because genetic structures act to restrict variation. If we look at phenotypic variation, we see that genetic heuristics preclude certain variations because they are fatal. Certain mutations, such as those which would lead to the absence of a heart in a mammal, are eliminated before the organism comes into being and so cannot show up as phenotypic variations. These genetic heuristics at the level of phenotypic variation restrict biological variation as well. In parallel to our broad definition of epistemic heuristics, genetic heuristics are the restrictions on biological variation. Following this line, we might say that biological evolution is built on some genetic heuristics which are, for example, chemical in nature. In other words, certain chemical structures need to evolve before biological evolution can begin. Thus, the disanalogy becomes an analogy: the innate heuristics in epistemic evolution are likened to genetic heuristics in biological evolution.

IV. OBJECTIONS

We will now go on to consider various objections to the explanation of the apparent guidedness of epistemic variation and to the strength of the analogy between biological and epistemic variation. Michael Ruse has argued that by attributing the guidedness of epistemic variations to epistemic preadaptation, evolutionary epistemology improperly locates the analogy in such a way that it breaks down. Ruse says that even if the

appeal to epistemic preadaptation were correct, evolutionary epistemology would “locate the supposedly Darwinian element in science at the wrong point to save the organism/science analogy, as it is usually conceived” (Ruse 1986, p. 59). More specifically, this objection sees evolutionary epistemology as arguing for a relation between biological and scientific entities, that is, beliefs about the world, or first-order beliefs. The move to epistemic preadaptation changes part of the relation to one between biological entities and heuristics, namely beliefs about beliefs or second-order beliefs.

This objection, we think, rests on an overly restrictive view of evolutionary epistemology as well as on a fuzzy distinction between first-order and “nth-order” beliefs. First, there is no reason for evolutionary epistemology to focus only on the relation between biological entities and first-order beliefs. Rather, evolutionary epistemology should be viewed as positing a relation between biological entities and beliefs at all levels (that is, first-order beliefs, second-order beliefs, etc.). On this view of evolutionary epistemology, the appeal to epistemic preadaptation does not locate the relation at the wrong level. Second, and we think more strikingly, there is no clear-cut distinction between first-order and second-order (or nth-order) beliefs. As Ruse (1986, p. 59) describes it, evolutionary epistemology deals with particular scientific discoveries, such as Watson and Crick’s model of DNA. Ruse argues that if we locate the analogy at the level of heuristics used by Watson and Crick, rather than at the level of their particular discovery, then we are improperly locating the analogy at the level of second-order beliefs. However, discoveries made by cognitive psychologists are beliefs about beliefs (second-order beliefs), and it seems that, in this case, Ruse would want the analogy to be located at the second level since it is at this level that the discoveries of cognitive psychologists are made. The point is that even if we accept Ruse’s suggestion to locate the analogy at the point of discovery, we are left with an open question as to which level of belief the analogy applies. Further, to return to two of the specific examples of discovery heuristics which we discussed above, both concrete heuristics and exemplars are instances of first-order beliefs that act as heuristics by restricting further epistemic variation.

Ruse, however, has a simple response. Rather than making a distinction by talking about levels of beliefs, he can do so by distinguishing between the subject of inquiry and the method of inquiry. Ruse could argue that evolutionary epistemology, through the move to epistemic preadaptation, places the analogy at the wrong place. Rather than locating epistemic variation at the subject of inquiry and comparing it to biological variation, by making the move to epistemic preadaptation, evolutionary epistemology locates epistemic variation at the level of both the subject and the method of inquiry. The adjusted version of Ruse’s argument is as follows:

1. Variation in biology takes place at the level of genetic mutation.
2. Evolutionary epistemology, in order to explain the seeming guidedness of epistemic evolution, needs to make the appeal to epistemic preadaptation.
3. The appeal to epistemic preadaptation locates epistemic variation at two levels: at the subject of inquiry and at the method of inquiry.
4. From 1 and 3, the analogy between epistemic and biological evolution breaks down because biological and epistemic variation are different.

We think this argument fails; it turns on dividing beliefs into two classes (first-order and *n*th-order), while leaving biology undivided. However, it is certainly possible to divide up biological variants in a manner similar to the way Ruse divides epistemic variants. We could talk of two classes of genes: regulator genes (that is, genes that affect other genes) and non-regulator genes (that is, all other genes). This division parallels the division between the subject of inquiry and the method of inquiry — in other words, the distinction between beliefs and beliefs about beliefs. Thus we could say that biological evolution also locates the variation at two levels and the analogy is preserved. If and how we divide biological and epistemic entities depends on our subjective viewpoint, that is, it depends on how we want to look at the matter. It seems that no philosophical mileage can be gotten from this sort of distinction.

Let us illustrate this point more clearly with another hypothetical example. Suppose I have lost my keys. The first place that I would look for them would be my coat pocket; the search is not a random one. Rather, my search is guided by the heuristic H_1 , which says: when I lose my keys, it is highly likely that I left them in my coat pocket. However, is H_1 really a heuristic or is it just a belief (a belief about my forgetful habits)? H_1 , in fact, is based on another heuristic, H_2 : when I lose something, it will probably be found in the place that I often leave it. H_2 is ultimately based on the inductive heuristic H_i : past experience is a guide to future experience. Where should we locate the analogy in this case? H_2 seems to be both a belief about the world and a heuristic, since it guides the variation that produces H_1 .

Another response to our adjusted version of Ruse's argument would be to attack premise 1 directly. Saying that variation in biology only takes place at the level of genetic mutation is to ignore the fact that there exist genetic heuristics, innate restrictions on biological variation. Genetic heuristics such as chemical preadaptations, for example, limit biological variation. The existence of heuristics at levels other than biological mutations is in direct opposition to premise 1. These genetic heuristics parallel the innate heuristics of epistemic variation and thus save the analogy. Thus, Ruse's argument, even if revised to make it stronger, does not seem to do any damage to the appeal to epistemic preadaptation.

Peter Skagestad argues that while the move to epistemic preadaptation may save the notion of the blindness of epistemic variation, it does so at the cost of undermining the central analogy of evolutionary epistemology. He argues that epistemic preadaptation “. . . may well act in a manner diametrically opposite to biological preadaptation” (Skagestad 1978, p. 615). He claims that a biological preadaptation increases the likelihood of further adaptations, while an epistemic preadaptation often decreases the likelihood of further adaptations by either leading its possessor to a conceptual dead end or by getting its possessor stuck in a rut.

This objection, we think, is based on a misunderstanding of both biological and epistemic preadaptation. Recall how preadaptation is used by biologists to explain, for example, how an organism develops a wing. The biologist’s answer involves the existence of a pre-wing structure, a half-wing, which had some preadaptive function. The move to biological preadaptation explains how a complex organ evolved. Evolutionary epistemologists use the move to epistemic preadaptation to explain the seeming guidedness of actions that they want to claim are the result of blind variation. Like the biologically preadapted structure, the epistemically preadapted heuristic had to have been good for something before its current adaptive use.

Skagestad’s objection is that a biological preadaptation increases the possibilities of adaptation, while an epistemic preadaptation reduces them. We think he is wrong in both cases. As we have already noted, both sorts of preadaptations create the potential for new adaptations as well as eliminating some of the previously existing potential for adaptations. For example, when birds developed half-wings, the possibility of evolving wings (and thus the possibility of flying) developed, but, the possibility of using the structure that preceded the half-wing (the quarter-wing?) for something else, say for picking up small sticks, disappeared (at least for the moment — the quarter-wing may reappear, though the chances of this are infinitely small). The same is true of epistemic preadaptations: when a scientist develops a certain heuristic, a set of possible discoveries that are based on that heuristic becomes available, but another set of previously possible discoveries is ruled out in light of the heuristic. There are dead end mutations in biology which sometimes get selected (just as there are heuristics which rule out good discoveries); we just do not usually call them failed preadaptations. In other words, while biological preadaptation suggests that the earlier organ enables variation, the presence of that organ just as surely restricts variation. This is particularly vivid when we consider animals that simply become extinct in a situation where different animals would probably have been able to evolve to meet the situation. The animals with the “right” preadaptations can adapt, while those with the “wrong” preadaptations die. Whatever we choose to call them, it is clear that in both biological and epistemic evolution there are variations which are selected but that are not ultimately adaptive. Although we have

emphasized their restrictive side, heuristics do make certain conjectures more likely. By eliminating some alternatives, heuristics have the effect of making the probability of other conjectures. In this way, heuristics, like biological preadaptations, both eliminate and create possible variations. Thus, biological and epistemic preadaptation act in the same manner and the analogy is preserved.

Paul Thagard (1980, pp. 188–190) and William Bechtel (1984, p. 316) argue that while biological variations arise truly independently of their environment, epistemic variations do not. As Michael Bradie writes, “. . . conjectures and theories, although they may be blind with respect to their potential success, are not randomly generated with respect to the needs of knowers but are developed in response to a problematic environment” (Bradie 1986, p. 424). In other words, with epistemic variation, the source of the apparent design is environmental pressures; the data we have seen up to now influence our next conjecture. This difference between biological and epistemic evolution has been conceded by at least one evolutionary epistemologist. Steven Toulmin (1972, p. 337) has argued that this disanalogy is explained by the fact that in epistemic evolution, variation and selection are *coupled*, that is, that the selection factors influence the factors that cause the variation, while biological variation and selection are completely uncoupled since the occurrence of a genetic mutation is unaffected by the environment.

Toulmin does not see his abandonment of strictly blind variation for variation guided by a coupled relationship with selection as destroying the analogy between epistemic and biological evolution, because it preserves both epistemic and biological variation and selection. L. J. Cohen (1973, p. 48; 1974, p. 324) has argued that this is not enough; an evolutionary epistemology with coupled variation and selection is not a true evolutionary epistemology. We agree with Cohen in thinking that Toulmin drops uncoupled variation and selection far too quickly. At least as far back as William James, who expressed “no hesitation whatever in holding firm to the Darwinian distinction [of the uncoupledness of variation and selection]” (James 1880, p. 456) even with respect to what he called mental progress, people have denied that epistemic variations are influenced by their environment. James writes:

I can easily show that throughout the whole extent of those mental departments which are highest . . . the new conceptions, emotions, and active tendencies which evolve are originally produced in the shape of random images, fancies, accidental outbirths of spontaneous variation in the function activity of the excessively unstable human brain, which the outer environment simply confirms or refutes, adopts or rejects, preserves or destroys — selects in short, just as it selects morphological and social variations due to molecular accidents of an analogous sort. (James 1880, p. 446)

In fact, on our definition of heuristics, the coupling problem is nothing new. To see this, let us return for a moment to our discussion of heuristics

and preadaptation. The appeal to epistemic preadaptation was made in order to explain the apparent guidedness of epistemic variation. The heuristics gained through preadaptation are just restrictions on variation: the more specific the heuristic, the more restricted the variation. Viewed in this way, the problem which coupling is meant to answer can be answered by the appeal to epistemic preadaptation. The data (the environment) influence the next conjecture (future variations); we do not need the coupling of variation and selection to explain this. Data influence the heuristics, which, in turn, restrict subsequent variation. This is the same in biology where the environment influences preadaptations, which, in turn, restrict future variations.

One further objection, made by both Thagard (1980, p. 190) and David Hull (1982, p. 307) is that the agents involved in epistemic variation want to solve a problem or accomplish a goal while the agents of biological selection want no such things. Even if the desire of epistemic agents does not guide variation, the mere fact that there is this desire seems to create a disanalogy: a scientist *wants* to solve a problem, even if her wants do not restrict variation. This is very different in biological evolution, where the members of the evolving species do not generally have any desires at all. As Thagard says, "The relevant difference between genes and theories is that theories have people trying to make them better" (Thagard 1980, p. 190). Because epistemic variation involves intentions, while biological variation does not, the analogy seems to break down.

Our response to this objection again involves the notion of the epistemic preadaptation of heuristics. Does the fact that epistemic variation involves intentions actually affect variation in any substantive way? If it has no effect, then it does not cause a disanalogy. If it has an effect, the only effect it could have is to restrict variation, in which case intentions can be thought of as heuristics, for heuristics are nothing except restrictions on variation. However, if these features can be thought of as heuristics, then they can be explained as preadaptations. In the face of all of these objections, we think that our resolution of the anomaly of guided variation survives.

V. CONCLUSION

This essay was motivated by the claim that there is a fundamental disanalogy between biological and epistemic evolution: biological variation is blind, while epistemic evolution is not. The problems an organism faces do not influence the new genotypes that will occur, but the problems a scientist faces seem obviously to influence the hypotheses she will propose. We have argued that this objection to evolutionary epistemology

is misguided. Some epistemic variation is unguided: some of our candidate hypotheses are in fact produced quite blindly, though this chaos is latter hidden for a variety of reasons. In other cases, epistemic variation is restricted, but by a mechanism that is analogous to biological preadaptation. In biological evolution, complex organs evolve from other organs. The earlier preadaptations make complex organs possible. They also effectively restrict biological variation, since the form that a complex organ can take is strongly influenced by the particular preadaptations available. Similarly, we have argued that epistemic variation is often restricted by preadaptations in the form of heuristics, and we have suggested what some of these heuristics are like. We have gone on to discuss the sources of these heuristics and have argued that, while some of them are presumably innate and so beyond the ken of a properly construed evolutionary epistemology, most of them are ultimately based on blind variation; that is, they are heuristics retained from previous variation.

Darwin's theory explains the surprising fit between organism and environment in terms of hidden randomness and selective retention; evolutionary epistemology uses the same sort of mechanism to explain the surprising fit between beliefs and the world. In both cases, we achieve an elegant and deep explanation of the way order can arise out of chaos. In both cases, we must also face the anomaly of apparently guided variation — the problem of complex organs in biology and of intelligent conjectures in epistemology. What we have argued in this essay is that both anomalies can be solved by an appeal to preadaptation, the application of previous adaptations, themselves the result of chance and inheritance. The model of random variation and selective retention is saved by showing that what appears to be non-random variation is in fact the result of selective retention. Some conjectures really are guided, but the guides are simply parts of the inheritance that is an essential feature of the evolutionary model. We have thus tried to save the central claim of evolutionary epistemology while at the same time preserving the strong intuition that the process by which scientists generate their hypotheses is not usually a random walk. We also hope to have provided some indication of the fruitfulness of fleshing out the analogy between epistemic and biological evolution.

NOTES

* We would like to thank Paul Bloom, Donald Campbell, Philip Clayton, Paul Pietroski, Peter Skagestad, Nicholas Thompson, G. L. Vankin, Laszlo Versenyi and two anonymous referees for their helpful comments.

¹ Our distinction between offering a biological explanation of belief and using biology as a model for a non-biological explanation parallels Michael Bradie's distinction between the

evolution of cognitive mechanisms program (EEM) and the evolution of theories program (EET) (Bradie 1986, pp. 403–413). In this terminology, our interest is in EET's. As Bradie observes, however, the relationship between these two programs is complex, since many of those who pursue the model approach also appeal to biology directly. That is, they include some EEM inside their EET. Our version of EET will avoid this complication.

² The idea that the guidedness of epistemic preadaptations can be explained by something analogous to biological preadaptations is discussed by Campbell (1974a, 1974b) and Skagestad (1978).

³ Gould and Lewontin (1978), among others, have extensively criticized what they call the adaptationist program. In particular, Gould has argued that 'preadaptation' is a bad word for the phenomenon that we use it to refer to because it implies that the new use of the structure was in some way anticipated. Instead, Gould recommends the use of the word 'exaptation'. We will use the more standard 'preadaptation', but it should be read as synonymous to Gould's term, i.e., *not* as implying any foresight.

⁴ Setting up this parallel is not to gloss over the differences between pre-Darwinian biology and pre-evolutionary epistemology. The pre-Darwinian view believes in perfect fit and no variation. The pre-evolutionary epistemology view believes in some variation and no *blind* variation. The pre-evolutionary epistemology view of epistemology is much like a sophisticated teleological argument in biology which holds that there is variation, but that this variation is guided.

⁵ Campbell (1974a) presents a variety of "testimonials" which give accounts of some of the ways that the blindness of epistemic variation is hidden.

⁶ There is a large literature on this reduction relation between new and old theories. One classic discussion is Ernest Nagel (1961), chapter 11. For a criticism of the traditional view, see Kuhn (1970) especially chapter IX. Kuhn's stand on incommensurability leads him to deny that new theories entail their predecessors as approximations, but he does seem to allow for concrete rules. Thus, he says that "[s]ince new paradigms are born from old ones, they ordinarily incorporate much of the vocabulary and apparatus, both conceptual and manipulative, that the traditional paradigm had previously employed" (Kuhn 1970, p. 149).

⁷ Kuhn maintains that those periods of normal science where exemplars serve as heuristics are not ones where theories are being tested. This may seem to disqualify exemplars as discovery heuristics, since these heuristics are supposed to generate hypotheses to be tested. But this may not be a serious problem. First of all, Kuhn has probably exaggerated the extent to which the normal scientist accepts her theoretical framework without question. More importantly, even if the over-arching theories the scientist uses to solve her problems are not up for testing, the particular solutions she proposes certainly are, and these solutions are themselves hypotheses.

⁸ We use "belief" as a broad term meant to include accepted hypotheses and accepted restrictions on epistemic variation; we do not, however, mean to count candidate hypotheses as beliefs.

⁹ Peter Skagestad (1978, pp. 615–616) defends such a view which he attributes to Charles Sanders Peirce.

¹⁰ Another empirical question is that of those theories not generated by innate heuristics, how many are generated by acquired heuristics and how many rest on hidden chaos?

¹¹ There are presumably possible worlds in which creatures very much like us in all other respects have all their heuristics and beliefs genetically hard-wired in them. The arguments that follow support our intuitions that this world is not such a world.

¹² It could be that the foundation of new exemplars is guided by some inter-paradigm heuristics, but Kuhn's emphasis is rather on the absence of such things. Kuhn's point is that if you overthrow the dominant paradigm together with its exemplars, what is left (e.g., inductive and deductive logic) radically underdetermines the new paradigm.

BIBLIOGRAPHY

- Bechtel, William: 1984, 'The Evolution of Our Understanding of a Cell: A Study in the Dynamics of Scientific Progress,' *Studies in History and Philosophy of Science* **15**, 309–356.
- Bradie, Michael: 1986, 'Assessing Evolutionary Epistemology,' *Biology and Philosophy* **1**, 401–459.
- Campbell, Donald: 1974a, 'Evolutionary Epistemology,' in P. A. Schilpp (ed.), *The Philosophy of Karl Popper*, vol. 1, The Open Court Publishing Co., LaSalle, IL, pp. 413–463.
- Campbell, Donald: 1974b, 'Unjustified Variation and Selective Retention in Scientific Discovery,' in F. J. Ayala and T. Dobzhansky (eds.), *Studies in the Philosophy of Biology*, Macmillan, London, pp. 139–161.
- Campbell, Donald: 1979, 'Descriptive Epistemology,' unpublished manuscript, William James lectures, Harvard University. Forthcoming in E. S. Overman, (ed.), *Methodology and Epistemology for the Social Sciences: Selected Papers of Donald Campbell*, University of Chicago Press, Chicago.
- Chomsky, Noam: 1986, *Knowledge of Language: Its Nature, Origins, and Use*, Praeger Publishers, New York.
- Cohen, L. J.: 1973, 'Is the Progress of Science Evolutionary?,' *British Journal for the Philosophy of Science* **24**, 41–61.
- Cohen, L. J.: 1974, 'Professor Hull and the Evolution of Science,' *British Journal for the Philosophy of Science* **25**, 334–336.
- Dennett, Daniel: 1984, *Elbow Room: The Varieties of Free Will Worth Wanting*, MIT Press, Cambridge, MA.
- Gould, Steven Jay and Richard Lewontin: 1978, 'The Spandrels of San Marcos and the Panglossian Paradigm: A Critique of the Adaptationist Paradigm,' *Proceedings of the Royal Society of London B* **205**, 581–598.
- Hull, David: 1982, 'The Naked Meme,' in H. C. Plotkin, (ed.), *Learning, Development and Culture*, John Wiley and Sons, Ltd., New York, 273–327.
- James, William: 1880, 'Great Men, Great Thoughts and the Environment,' *Atlantic Monthly* **46**, 441–459.
- Kuhn, Thomas: 1970, *The Structure of Scientific Revolutions*, second edition, University of Chicago Press, Chicago.
- Nagel, Ernest: 1961, *The Structure of Science*, Routledge and Kegan Paul, London.
- Nozick, Robert: 1974, *Anarchy, State and Utopia*, Basic Books, Inc., New York.
- Ruse, Michael: 1986, *Taking Darwin Seriously*, Basil Blackwell, Oxford.
- Skagestad, Peter: 1978, 'Taking Evolution Seriously: Critical Comment's on D. T. Campbell's Evolutionary Epistemology,' *Monist* **61**, 611–621.
- Thagard, Paul: 1980, 'Against Evolutionary Epistemology,' in Asquith, P. D. and Giere, R. N., (eds.), *PSA 1980*, 187–196.
- Toulmin, Stephen: 1972, *Human Understanding*, vol. 1, Oxford University Press, Oxford.