

William Harms  
11/25/98  
Draft Version Only

**PRE-ESTABLISHED HARMONY:**  
*Information Flow in Evolutionary Processes*

Introduction: toward a *third-person epistemology*

Part I: Generalizing Evolution

**Chapter 1:** Replicator Theories

**Chapter 2:** Cultural Transmission

[ “Cultural Evolution and the Variable Phenotype” *Biology and Philosophy* ...]

Part II: Modeling

**Chapter 3:** Single Populations

**Chapter 4:** Multiple Populations

**Chapter 5:** Information [ “The Use of Information Theory in Epistemology”, *Philosophy of Science* ...]

**Chapter 6:** A two level model for Bacterial Epistemology

**Chapter 7:** A three level model for Bumblebee Epistemology [ “Reliability and Novelty”,  
*Erkenntnis* ...]

Part III: The Question of Normativity

**Chapter 8:** Teleonomic Agency

**Chapter 9:** Is and Ought

Appendix: *A proof of information gain in selection driven processes*

References

Bibliography

## INTRODUCTION

There are a couple of different things one might be after in attempting a constructive contribution to epistemology. Descartes, for instance, was clearly after what I will call a “first-person” epistemology — an objective made clear by the very title of his most often read *Meditations on First Philosophy*. Descartes attempted to prove that *I* exist, and that from *my* existence *I* can infer certain facts about the manner of *my* creation, which in turn imply certain consequences regarding the reliability of *my* senses and mind. The payoff, if Descartes’ self-evident inferences compel *me*, is *my* being justified in accepting the evidence of *my* senses in the situations where *I* have deduced them to be reliable. Unfortunately for Descartes’ projects, not all of the *I*’s out there shared his sense of the self-evident. But if the project had worked, it would have offered an epistemology that could be *used* in epistemic deliberations to justify beliefs, a first-person epistemology.

Epistemology in the early twentieth century eschewed Descartes’ introspective methods as well as his appeal to God as epistemic guarantor, instead building on the empiricist doctrine that all knowledge comes via perception. Consequently, science, if it was to be justified, was to be reducible via the self-evident resources of logic and set theory (less self-evident) to observations. Again, if this project worked, one would have an epistemology that could be used in epistemic deliberations to justify beliefs. If one knew just which observations constituted the verificationist meaning of “electron,” for instance, then one would could use this knowledge to justify belief that an electron was present, doing whatever it is that electrons do.

Like many great ideas, this one sounded good until the details were worked out. One consequence of the development of this theory, principally through the work of Carnap, was the gradual weakening of its justificatory implications. In the end, as Quine convinced just about everybody in his (1969a) “Epistemology Naturalized”, the old logical empiricist epistemology was dead in the water. Quine, perhaps the quintessential late twentieth century empiricist, concluded that,

Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies natural phenomenon, viz., a physical human subject. This human subject is accorded a certain experimentally controlled input — certain patterns of irradiation in assorted frequencies, for instance — and in the fullness of time the subject delivers as output a description of the three dimensional external world and its history. (1969a, 83)

---

Psychology, then, because it was the laboratory setups of the cognitive scientist that were to be the primary tools of inquiry into human knowledge. Quine's assault on normative, which is to say first-person, epistemology was complemented by arguments by Ludwig Wittgenstein and followers to the effect that the illusion of epistemological problems is based on a misunderstanding of language. Richard Rorty (1979) influentially pressed this case, and there are many who follow him in believing that epistemology is today "an unnecessary discipline." Meanwhile, Edmund Gettier's (1963) formula for causing trouble for definitions of knowledge added to the impression that traditional epistemology was on the rocks.

Of course, there are still epistemologists, even first-person epistemologists. For if one concludes that empiricism leads to a naturalism that must abandon the traditional first-person emphasis of epistemology, then one is free to retort, "so much the worse for empiricism" (and, by implication, for naturalism as well). There are theoretical resources beyond those sanctioned by the logical empiricist austerity program for theory building, and one need not reject science in order to engage in productive discourse in terms that science cannot account for. Besides, if Quine's naturalized, or "third-person" epistemology is just part of psychology, it is not at all clear what business philosophers have with it, other than usual second-order study of scientific disciplines. As a consequence, a survey of academic philosophers might give the impression that "epistemology naturalized" is just as dead in the water as the old logical empiricist program.

Such an impression would be mistaken, however. For despite the rarity of anyone proclaiming themselves a "naturalistic epistemologist," there is plenty of interesting theoretical work being done by philosophers which bears on the project of *third-person* or descriptive epistemology. It's just that such contributions tend to fall under other categories, such as "philosophy of mind", "philosophy of language", or "decision theory." And traditional epistemologists, frequently wearing the double hat of "epistemologist/philosopher of science," are free to be intrigued by the ongoing developments in cognitive science without committing themselves to pulling an oar in Quine's boat.

*Evolutionary* epistemology, which this book hopes to advance, is an approach which explicitly styles itself as a naturalistic, third- person epistemology, though different in important

---

respects from what Quine seems to have had in mind. Evolutionary epistemology aims at a third-person theory, in that its primary aim is not to provide epistemic deliberators with tools for justification, but to provide a scientifically adequate and interesting analysis of knowledge systems. Indeed, the hope is to build an epistemology which can both use and contribute to ongoing research in evolution and cognitive science. To pursue such a project is not to insist that the traditional epistemological project is impossible, however, nor is it to insist that the descriptive approach will have no normative implications. Rather, it is an expression of the conviction that headway in the understanding of knowledge may be more easily made by deferring (or ignoring) first-person questions of justification until an adequate understanding of the systems involved is available. The issue of what if any justificatory implications the resulting theory will have can be best assessed once that theory is in hand. The commitment to naturalism just means that one approaches human knowledge as a natural (biological) phenomenon, to be understood in relation to simpler systems in simpler creatures. It also means that one can help oneself to the resources of natural science in building the theory, whether it be cognitive science, evolutionary biology, or what have you. Inevitably, traditional disciplinary boundaries become blurred, just as the boundaries between decision theory, economics, and evolutionary theory have been. As for philosophical reservations, Quine noted that the appearance of circularity involved in attempting to (ultimately) explain science in terms of science is not to be feared once one realizes that what is aimed for is not a foundational epistemology, but rather, one that coheres in the required way with the rest of our body of scientific knowledge, and performs well in that context.

The important difference between evolutionary epistemology and Quine's vision of epistemology as a chapter of psychology is the extent to which the former insists on the centrality of Darwinian evolution to understanding human knowledge. Psychologist Donald T. Campbell is primarily responsible for the term "evolutionary epistemology", and for *its* central vision. In his (1974) tribute to the late Karl Popper, whose later work anticipated the current project in important ways, Campbell laid the groundwork for later efforts, including the present one. Campbell believed that the Darwinian formula of adaptive change by variation and selection, understood in appropriately broad terms, was in fact not confined to the biological evolution of genes, but was instead ubiquitous. Not only genes, but also beliefs, scientific theories, languages,

---

the navigational behaviors of protozoa, and a host of other natural and social processes evolve in a Darwinian way. A central idea was that the results of selection on one level (like the genetic evolution of perceptual and sensory mechanisms) can act as selection mechanisms at other, higher levels (like trial-and-error learning). It was in the interrelation between evolutionary processes in a “nested hierarchy of blind variation and selective retention processes” that the key to understanding how human knowledge connects to the world was to be found.

One of the most common reasons for resistance to understanding human behavior in biological terms is the extent to which human behavior is variable under cultural influence. This familiar fact about human behavioral flexibility has no doubt been a major contributor to the long-standing tendency to treat the human mind as *tabula rasa* — a blank slate over whose content experience and culture had complete control. More recently, the rather extreme negative reactions to E.O. Wilson’s *Sociobiology* (1975) and *On Human Nature* (1978) were to some extent justified (if not motivated) by the difficulty of understanding the relevance of innate structures in a creature for whom any of those structures can apparently be overridden by culture. The current trend in evolutionary psychology (e.g., Barkow, Tooby and Cosmides, 1992) continues to suffer from skepticism on the same grounds.

It might seem as though evolutionary epistemology faces the same difficulty rather inevitably. But here, one can begin to see the importance of Campbell’s notion of a selection hierarchy. For if genetically inherited mechanisms act as *selection* mechanisms on essentially Darwinian processes of behavior acquisition and modification, then one would expect to see pretty much what one does see — relatively few constraints on what one *can* do, combined with systematic tendencies to continue or cease new behaviors based on the response of the inherited mechanisms. The simplest cases relevant to human behavior involve inherited pain and pleasure sensors, whose functioning inhibits or reinforces acquired behavior. One can try just about anything, but one will tend to stop doing the things that hurt, and continue doing what brings pleasure. This basic idea, extended to accommodate multiple genetically determined mechanisms and selection hierarchies with more levels, provides a rich framework within which to study genetically determined constraints on human behavior, a framework which is consistent with the observed variability of that behavior.

---

A second predominate reason for skepticism regarding the potential of naturalism rests on a feature of empiricism itself. Traditionally, empiricism has focused on *occurring* causal interactions (or at best current dispositions) to the exclusion of all else. This exclusive focus played centrally in Hume's analysis of the is/ought distinction.<sup>1</sup> And Quine, again in "Epistemology Naturalized", insisted that for the empiricist, "all inculcation of the meaning of words rests ultimately on sensory evidence" (75), a point he made repeatedly. It is only in the last fifteen years or so, principally through the work of Ruth Millikan and Fred Dretske, that the relevance of the essentially historical relationships of evolutionary biology to meaning has been discussed. Though "evolutionary semantics" remains a minority view, it is no longer possible to insist with Quine that naturalistic meaning must depend purely on actual stimuli. Evolutionary theory, unlike the physics which constituted the paradigm of science for the early twentieth century empiricists, turns crucially on historical relationships of common ancestry and adaptive functioning in historical environments. Traditional empiricism hobbled naturalism by assuming that the only facts relevant to understanding human behavior were those concerning occurrent causal interactions. Evolutionary theory provides new resources which change the game.

The tendency of empiricism to focus on occurrent relationships has resulted in empiricist accounts of norms which seem to reduce them to expressions of emotion, there being apparently no resources for understanding "rules" which are to be followed other than those dictated by capricious culture, or those which constitute law-like regularities of natural phenomena. As naturalistic meaning theorists have shown, however, the rules involved in the evolutionary design process have a different character than either of those sorts recognizable by empiricism. They are more objective than the dictates of culture, being entrenched as they are in history, and yet it is possible not to follow them, unlike laws of nature. Consequently, there seems to be the potential for building a naturalistic theory of norms which is richer than those possible under empiricism, and which does greater justice to the phenomenon itself.

---

<sup>1</sup> This will be discussed in Chapter 9.

The situation, then, seems to be that evolutionary theory provides resources for the “naturalization” of epistemology far beyond what Quine had in mind, both in the form of models where cultural evolution interacts with biological evolution, and in the form of the rules of adapted functioning and other essentially historical resources which promise a more satisfying third-person account of norms and meaning. Yet despite significant contributions to the project by philosophers like Clifford Hooker (1995) Aharon Kantorovich (1993), and relevant research by a host of others, evolutionary epistemology itself has yet to find focus as a contending approach to the understanding of knowledge. There are any number of reasons for this. This book tries to address three of them.

Part I (Chapters 1 and 2) deals with what one might think of as the ontology of cultural evolution. Richard Dawkins (1976), in the process of trying to cure people of “group selectionism” in biology, convinced a lot of us that, strictly speaking, evolutionary processes are *copying* processes, and that the essential ingredient for any process to be properly Darwinian is for it to consist of the dynamics of lineages of self-replicating entities, or “replicators.” David Hull (1988) ran with this idea in his analysis of scientific progress, as has Daniel Dennett (1991, 1995) in various pursuits in the philosophy of mind. The result is that many believe that, in order to have a theory of cultural evolution, we need to buy into the existence of “virus-like informational entities” that “leap from head to head.” In Chapter 1, I analyze the notion of a replicator in some depth, both in Dawkins’ and Hull’s formulation, and conclude that the general existence of cultural replicators, or “memes”, is problematic. Any squeamishness with respect to a theory of cultural evolution which rests on them seems to be well justified. In Chapter 2, I present an alternative way of thinking about cultural transmission which does not require replication, characterizing it as coordinated phenotypic variability among closely related creatures. This opens the door to a more abstract characterization of evolutionary processes in Part II.

It is curious that, despite the lack of a generally available evolutionary account of *moral* norms, “evolutionary ethics” is a recognizable philosophical activity to a degree that evolutionary epistemology is not. This, I think, is due to the existence of applications of evolutionary game theory which provide a rigorous framework within which the evolution of phenomena like

---

altruism can be investigated. Perhaps most influential in this respect is the work of Robert Axelrod (1984) in promoting the evolutionary analysis of cooperation in the “prisoner’s dilemma.” Many have followed, and today the performance of “TIT-FOR-TAT” in evolutionary models is familiar to those in a multitude of disciplines, where systematic progress continues to be made. One problem with evolutionary epistemology, it seems, is that it is not at all clear what an evolutionary epistemologist *does*. Perhaps what is needed is a framework within which one could *build* models of something like Campbell’s selection hierarchies and *study* the behavior of multi-level selection systems.

Part II (Chapters 3-7) is devoted to the construction of such a framework. Chapter 3 develops the conceptual basis for an appropriately general notion of evolutionary processes and the associated formal basis for modeling them. The formalism is a variant on the population models that form the core of population genetics and evolutionary game theory. Chapter 4 extends the framework to accommodate multiple populations and their interactions into a “network” of evolutionary processes capable of modeling the flexible constraints on culture suggested by Campbell. Chapter 5 borrows the “mutual information” measure from information theory for the purpose of assessing the efficiency with which populations of acquired traits track the environment. This use of information theory is compared to that of Dretske (1981), and some preliminary results regarding the relationship between information and fitness are established. A proof of information gain under selection is contained in the appendix.

Chapters 6 and 7 are devoted to simple two and three-level selection models which illustrate basic principles operating in multi-level selection systems. Chapter 6 explores the basic “fitness feedback loop” in a two-level model of bacterial navigation, which maintains exploitable information in distributions of acquired traits via the interaction between genetic selection on sensor configurations and selection in the bacteria’s random walk. Chapter 7 extends this model to a three-level model of bumblebee foraging. The introduction of middle-level preference formation demonstrates how learning about novel items can emerge through the exaptation of success indicators, and it turns out that information caught in middle-level distributions has suggestive properties for epistemology.

Even if what one is after is a third-person descriptive theory of knowledge, building on

---

and contributing to theoretical work in cognitive science and evolutionary biology, the inability of empiricism to account adequately for norms may still seem too high a price to pay for a naturalistic epistemology. The eliminativist argument that, as science progresses, bits of mythological vocabulary must fall by the wayside is well taken, perhaps. Just because it was once convenient to understand the operations of nature in terms of various occult forces does not mean that science must forever take the existence of these forces seriously. Nonetheless, to think that anything that empiricism cannot account for may be safely disregarded would be going too far — especially since we are just beginning to understand the relevance of adaptive histories to normative objectivity.

Part III (Chapters 8 and 9) addresses the adequacy of a third-person evolutionary account of norms. Chapter 8 presents central concepts from Ruth Millikan's theory of "proper functions", though it is argued that naturalists do not require *definitions* of terms such as "function" where their use is sanctioned by evolutionary biology. It is shown how many of the familiar features of *agency* are present in complicated adapted systems. Chapter 9 shows how a basic understanding of the semantics of evolved signaling systems illuminates the is/ought distinction, and concludes by returning to the question of epistemic justification. It is concluded that, at least in principle, evolutionary naturalism *can* identify objective truth conditions for "is justified" or "doesn't make sense", though this falls just short of being able to *actually* justify beliefs.

The aim of this book, then, is to construct tools which will help us to *do* evolutionary epistemology. This should, of course, be of interest to philosophers attracted to naturalism. But it seems to me that the utility of the tools involved extends to fields outside of philosophy. In particular, multi-level selection models provide an alternative to genetically determined "programs" in modeling flexible constraints on culture and learned behavior in general. This broadens the resources available to sociobiologists and evolutionary psychologists for understanding how evolution affects human behavior. Similarly, the functional analysis of adapted systems allows richer accounts of meaning and norms than those common in cognitive science, yet which obey the third-person descriptive scruples of modern science. This opens the door to productive interchanges between evolutionary theory, cognitive science, decision theory and other

traditionally normative fields in philosophy.

This is not to say, however, that epistemology “naturalized” is to become merely a chapter of psychology, or of any other scientific discipline. The resources required span so many disciplines that the evolutionary epistemologist must of necessity be a new kind of specialist, one with sensitivity to the philosophical issues and facility with a variety of analytical tools. In the end, I do not expect that a traditional first-person epistemology will be forthcoming from the project. On the other hand, we may well be surprised by the extent to which evolutionary naturalism *can* address the traditional issues. But as for that, only time will tell.

*San Diego, October 1998*