Information and Meaning in Evolutionary Processes

William F. Harms. Ph.D.

Draft: 12-1-1

[Forthcoming from Cambridge University Press]

Introduction:

Why Epistemology Matters

The need to understand knowledge itself is felt by people in different ages for different reasons. The justification of religious truths and ethical precepts, the desire for certainty in an uncertain world — both have been effective at motivating inquiry into the fundamental nature of knowledge. Our own need can, perhaps, be best illustrated by a story. The story I have in mind involves a large rat trap. It also involves two young university professors, a celebrated theory war, about eight hundred undergraduate students, and a book by Thomas Kuhn. It is a true story, or at least, it started out as a true story. This is how I remember it.

Young Professor B and young Professor S had been recruited from their various departments, Philosophy and German Literature, respectively, to participate in an interdisciplinary extravaganza known as "Core Course." Core Course was the sort of politically correct Great Books program common in large universities. Drudge work for the 800 students (assignments, grading, one-on-one contact, and daily class time) was handled by 20 or 30 "teaching associates" (second year graduate students), staff meetings once a week. Young Professors S and B were accorded starring roles in this production, giving primary content lectures to the eight hundred students in two shifts, the auditorium being limited to half that number. In their home departments, Professor B's job was to do something called "philosophy of physics" which involved a lot of mathematics, and to teach undergraduates things like logic and critical thinking. Professor S's job was to interpret great works of German literature, and to teach students about that and also about things like postmodernism and deconstruction. Needless to say, Professors B and S disagreed about some very basic things. This was the Winter quarter session, which included material on the History and Philosophy (and Sociology) of Science.

The incident of interest occurred during the lectures on Thomas Kuhn's very influential *The Structure of Scientific Revolutions*. In case you've forgotten, this is the book that got everyone talking about "paradigm shifts" in the 1960's. Kuhn's basic thesis was that scientific progress cannot ever be the steady accumulation of truths over centuries that its proponents sometimes claim. Instead, science progresses through a series of lurching cycles driven by the acquisition of various bits of scientific infrastructure — concepts, measurement apparatuses and model studies — followed by the discovery of the limits of that technology. This results in a cycle of crisis and revolution in which the old infrastructure is replaced by new, frequently accompanied by an actual loss in explanatory power for the discipline in question. Critics of science had a field day with this stuff, much to Kuhn's dismay, claiming that he had shown that scientific objectivity is a mere social and institutional fiction, science having no more inherent value than witchcraft or performance art. This is, of course, not what Kuhn had in mind, but that's a different story.

Professor S had first go at this material, and took the opportunity to promote some of *his* favorite causes — that reality is a social construction, that political power drives the content of all belief systems, and that while the science boys seem to have been getting a lot of the attention (not to mention money) lately, what this means is simply that they have the power, no more. Lore is lore. The Upanishads and Newton's *Principia* are on a par in this respect.

Young Professor B could not take this lying down. So when his turn to lecture came, he showed up with the aforementioned rat trap, his intent being to challenge Professor B to a sort of practical duel of ideologies. The lethal efficiency of the rat trap was demonstrated — lethal to a pencil anyway — and Professor S was offered the opportunity to demonstrate the much vaunted social construction of reality. Surely, if Professor S truly believed reality to be a construction, then he should be quite happy to put his finger in the trap, secure in the knowledge that the trap could not *really* hurt him if he believed it would not. Quick on his feet, Professor S replied that he was sufficiently at the mercy of the way his own society constructs reality that he was quite sure that the trap would actually hurt him. But, he insisted, someone else from a society which constructed reality in a sufficiently different way might not be harmed. And with this impasse, it was time for the bell.

To anyone who is even marginally aware of what has been going on in the Humanities in major U.S. universities in the last forty years, the theoretical conflict that lays behind this incident is familiar, and if anything, the incident itself was relatively free from the acrimony one often finds in less public conversations on the matter. Analytic philosophers can be viciously critical of the "fuzzy headed thinking" in so-called continental philosophy, and of the destructive potential of unchecked relativism, historicism, and deconstruction. Continental philosophy, usually a bigger force in the Humanities at large than in philosophy departments proper, views

analytic philosophers as insufficiently aware of the social, historical, and political contexts in which humans live and of the effect these contexts have on how people understand themselves and their world. The left claims that the right is in denial about the constructedness of reality. The right claims that the left has given up its right to a place in an institution dedicated to the furtherance of human understanding.

As is usual in these kinds of disagreements, both sides have a point. On the one hand, human knowledge is not discovered but constructed. The beliefs we have are constructed from the concepts available. Those concepts are in turn constructs, and while one might be able to argue that for a given purpose some set of concepts is optimal, there seems little reason to think that some one set of concepts is optimal for all purposes of all species at all times. On the other hand, even accepting the constructedness of our conception of reality, one may reasonably insist that not all such constructions are on a par from a practical point of view, and that human beings and the things they do are sufficiently similar that some world-views might very well be flat out better than others.

Ideally, one would think that this crucial philosophical question about the status of knowledge would be the proper subject matter of epistemologists, and that epistemology as a discipline would step up to clarify these matters. But this has not been the case. Outside of analytic philosophy many believe, as Richard Rorty (1979) argued, that the death of foundationalism has left epistemology as an impossible and unnecessary discipline, given that epistemologists have traditionally attempted to discover some area of human belief that transcends the possibility of doubt. Meanwhile, epistemology itself seems to have become something that philosophers do "on the side" when epistemological issues arise in their areas of primary interest, or something done by specialists who work on small areas of the large epistemological puzzle. Looming is the question of how, in the face of the constructedness of our knowledge of reality, to say anything non-question-begging about the relationship between that conception and the world we want to believe it mirrors.

Why Evolution Matters

What, if our conception of reality is a construct, can epistemology possibly salvage? Two possibilities come to the fore. First, we ordinarily have a certain amount of confidence that the care we take in acquiring and evaluating beliefs is not a complete waste of time. It generally seems to us that beliefs have something to do with the reality "out there," and ordinary perception and common sense are largely to be trusted. What reason can we have to trust our senses, minds, and memories? What reason can we have to think that our cognitive and perceptual apparatus is reliable? Such questions suggest themselves for our consideration even if we accept that certainty is not to be had, and even if we acknowledge that there is probably no such thing as *the* truth for all people at all times.

Second, "knowledge" itself is an evaluative, normatively loaded term, and our conception of knowledge is replete with a number of such concepts. Truth, justification and meaning all demand explanations. The very concept of knowledge implies more than just that our mental states reliably and usefully track the world. It implies that there is such thing as getting it right, that there are some sort of *rules* which apply to how we form and interpret beliefs that go beyond mere usefulness. Moreover, it is not only the rules regarding knowledge that we claim to know. We claim to know the difference between right and wrong, just and unjust. We claim that there is something to *know* about such matters, but what might ground such knowledge is more of a mystery than the basic question of how our concepts and beliefs relate to the world. At least in the latter case, we can ask how concepts relate to whatever there is "out there" that makes things happen. In the case of the normative question, the various sorts of "oughts" or rules do not even seem to be a matter of what is going on with the "hidden springs and principles" of nature, to borrow Hume's phrase. The purported rules are a matter not of what is, but of how things ought to be.

Together, these two questions concerning reliability and normativity present the basic challenge for epistemology. In this book I will try to show that, in both cases, the answers come from understanding evolutionary processes.

In the first case, the general solution is not so hard to see, nor is the suggestion particularly novel. Briefly, the only reason to think that our thoughts usefully and reliably mirror reality is that if they didn't at least get us coping fairly well with the world, it's not likely we would be here at all. The principle we are alluding to is *natural selection*, the driving force of biological evolution, and the school of thought that emphasizes the importance of natural selection for the understanding of knowledge is usually called *evolutionary epistemology* (Campbell, 1974).

Evolutionary epistemology differs from other sorts of naturalistic epistemology in the extent to which it emphasizes the importance of natural selection. "Naturalistic" epistemologies in general

eschew the traditional emphasis on introspection (as well as the traditional acceptance of at least the possibility that minds are more than mere matter) for the scientific point of view. For naturalists, human beings are material entities and anything that is of interest about them, including the phenomenon of knowledge itself, is to be studied empirically and scientifically. This in itself does not necessarily bring with it any commitment to the importance of evolution. On the contrary, classic naturalistic epistemologists like Quine (1969) insisted that epistemology was to consist of observing actual physical processes involved in perception, cognition, and the like, and while he noted that evolutionary history was responsible for the similarity of our quality space to the quality space of the world, reasoning about evolutionary history in any detail did not seem to be part of his program for the new epistemology. Evolutionary epistemologists, while also naturalists in their insistence on the status of human beings as material entities, believe that we cannot fully understand knowledge purely in terms of the currently observable causal processes that the physical sciences focus on. This is not to say that one needs any more than an understanding of current causal structures and processes to understand *how* we learn about our world. But understanding *why we are able* to so learn requires some understanding of our evolutionary history. Moreover, given that epistemology must ultimately attempt to give a coherent account of its own nature, understanding the general principles of how evolution adapts cognitive systems to environments seems essential to understanding the *nature* of the relationship between thoughts and the world. Even for those of us who have given up on foundationalism, some way of standing back and thinking about the fit between our own concepts and the world, as well as a way of understanding comparisons between different world views, is quite desirable. If we can understand something about how much and what kind of optimization of cognitive fit is likely to emerge from evolutionary histories, we will be that much closer to a more adequate scientific understanding of knowledge and more able to answer some of the interesting comparative questions about knowledge that arise between cultures and between species. The general idea of an evolutionary epistemology has been around for some time. The pressing need at the current juncture is for the development of theory and in particular formal tools of analysis that will make the relationship between selection and mind-world coordination more tractable.

As for the second question, evolutionary epistemologists have for the most part adopted the naturalistic party line in abandoning the traditional normative concerns of epistemology, concerns with capitalized topics like Truth, Justification, and Reason. Hume, very possibly the first modern naturalistic philosopher, argued that judgements of *moral* propriety were simply the operation of certain sentiments. Early twentieth century naturalism resulted in the similar thesis of emotivism, according to which normative value judgements are merely the expression of emotions, devoid of the kind of referential content that might make them objectively true or false. In this case, what applies to morality applies to Truth, Justification, and Reason as well. And indeed, the current majority opinion seems to be that naturalism unavoidably brings with it the re-analysis of any and all normative judgements in terms of the proximal emotions that give rise to them. The sting of abandoning so much of the traditional concern of epistemology may be salved to some degree by the thought that, after all, it does no good to complain that the things people in fact do are not the things they *ought* to do. On the contrary, the more energy we spend on understanding both how and why people do what they do the sooner we will be able to stop complaining and actually do something about human epistemic failings. Still, the most common reason for rejecting naturalistic epistemology is its alleged failure to accommodate the normative, and proponents of naturalistic epistemologies have had little to say in opposition (Kim, 1988). It would seem that either naturalism is incomplete, or the normative is to be ignored along with the rest of the supernatural. In either case, one won't be able to get an *ought* from an *is*. The current philosophical consensus seems to be that if both is (facts) and ought (norms) are to be studied they will be studied in fundamentally different ways, perhaps by very different disciplines. This consensus also maintains that all previous attempts to explain value on the basis of fact have failed. As a result, anyone attempting to account for the truth of normative judgements within a naturalistic framework has to contend not only with the inherent difficulty of the task, but with the fact that everyone already *knows* it can't be done, as well.

I shall argue that the problem with our current understanding of the gap between *is* and *ought*

is that it is based on a very outdated conception of the proper domain of scientific enquiry. Taking physics as the paradigm, we have conceived of science as concerned exclusively with occurrent causal processes. Like many such conceptions, this one has been subject to considerable erosion in the centuries since its emergence. Just as seventeenth century mechanistic materialism was forced to broaden itself to accommodate "action at a distance," first in the form of gravity and then in the form of electromagnetism, so the empiricist insistence that contingent history is merely anecdotal is gradually giving way to the recognition of the central importance to biology of essentially historical relationships like kinship and adaptation. The increasing importance of historical relationships is in turn driven by the increasing status of biology as a genuine science and the concomitant change in our conception of what constitutes a genuine science. It turns out that these historical relationships pretty much make a shambles of the arguments of David Hume and G.E. Moore which seem to lie behind the common "knowledge" of the intractability of the is/ought gap.

Exactly how the historical subject matter of evolutionary biology changes the big picture will be the subject of part III of this book. Roughly, I shall argue that a broad enough theory of meaning can account for the meaning and thus truth of both "is-statements" and "ought-statements," and in so doing account for the difference between them. In the absence of traditional assumptions regarding abstract meaning-entities (e.g., propositions), we are forced to recognize that meaning is conventional, and conventions are historical entities. This focus on the ways in which meaning conventions emerge highlights facts about meaning that are usually overlooked. In the end, the general analysis of the emergence of meaning conventions shows indicative language to be a rather specialized sort of signaling system, and not the only one capable of correspondence truth. Normative utterances are more akin to warning cries than statements of fact, untranslatable but not unanalyzable in our usual common mode of descriptive speech.

Given the general consensus that this sort of project cannot work, it behooves me from the start to be completely clear about my own aspirations with respect to a theory of the normative. I will argue that a full knowledge of the functional history of signaling systems of all sorts, including especially the system stabilizing consequences of signaling behavior, is sufficient to establish the conventions governing meaning and truth for all sorts of signals, from hormonal secretions to scientific hypotheses to pronouncements of moral and epistemic justification. The point, however, is purely academic in that the sorts of historical facts relevant to determining meaning-conventions are so difficult to come by that we should not ever expect to see our own intuitions on such matters overridden by pronouncements of evolutionary science. There is in addition a purely theoretical reason why you can't get an ought from an is, but we will be in position to actually explain this, rather than merely recognize it as we usually do. So even if the possession of all the relevant facts allowed us to establish the meaning of normative utterances and thus tell us whether they are true, this is not quite the same as *telling* us what we ought to do. As a consequence, both according to the strict letter of the theory and because of the unavailability of pertinent facts, the theory I will propose falls short of the traditional aspirations of epistemologists — to issue authoritative epistemic norms. I am enough of a naturalist that this doesn't bother me. What I do think is philosophically important is the way in which a defensible descriptive theory of the normative can counteract the relativism with respect to normative standards which has accompanied our materialistic world-view. If a theory can reassure us that, even in a purely material world, there may be standards of conduct and thought that are sufficiently general and objective to apply to *all* human beings, this is no small thing. Indeed, by my lights, we are far more in need of a theory which helps us make sense of the very possibility objective norms than we are of one which tells us what they are.

Ontology, Selection and Convention

The nine chapters of this book are organized into three parts. Parts I and II are concerned with furthering the general evolutionary epistemology project, Part I with ground-clearing, Part II with the construction of formal tools for analyzing multi-level selection and information transfer processes. Part III deals with the more contentious evolutionary take on meaning conventions for normative intuitions. The three parts can be read independently, and may appeal to different interests and temperaments.

Part I: Evolutionary epistemology, at least in its more ambitious versions, requires a way of thinking about evolution that is broad enough to allow it to

occur in culture as well as in the genetic lineage. The most common such conception is Richard Dawkins' (1976) "meme," a self-replicating "informational entity." Dawkins, along with a number of others, is convinced that evolution by natural selection only happens to lineages of self-replicating entities, and consequently, if evolutionary concepts are to apply to culture then there must be some sort of self-replicating entity in culture.

In chapter 1 I discuss at some length versions of the cultural replicator from Dawkins, David Hull, and Daniel Dennett. I conclude that there are not one but three different notions, and that none are adequate for epistemological purposes. In chapter 2, I broaden the critique to the nascent field of memetics in general, fielding alternatives and trying to set some standards for conceptual innovation. I argue that genes are not the only or even the best way of thinking about the tree of life — the lineage of dividing cells has a concrete identity over time that genetic "information" does not. Cells, moreover, may be the only thing that can truly be said to "self"-replicate, and the logic of that process is fission and regrowth, rather than transcription. Consequently, looking for gene-like entities as a prerequisite for cultural evolution is rather misguided. Nor does cultural transmission require an entity to be transmitted. On the contrary, communication and cultural transmission are easily understood as coordinated state change in closely related organisms. Finally, even if memetics is defended only as a way of looking at cultural evolution and transmission, there are limitations to its utility.

Part II: Chapter 3 begins my positive account, developing a general model of evolution derived from the formal models of population genetics and evolutionary game theory. The emphasis here is two-fold: First, on extracting general evolutionary concepts like selection, fitness and variation from our best abstract mathematical understanding of evolution rather than proposing causal analogs to the actual physical process of biological evolution as the replicator approach does. Second, on generating a simplified formal model suited for computer modeling of both biological and cultural evolution. Chapter 4 explores the mathematical concept of mutual information for the purposes of epistemology, and establishes some simple results regarding the utility of information. Chapter 5 puts the two together, establishing natural selection as an information transfer process. Chapter 6 develops a two-level selection and variation model interpreted as a model of bacterial navigation, with an eye to creating a formal model for the basic interdependencies between biological and cultural evolution. Chapter 7 develops a three-level model of bumblebee foraging, which accommodates the formation of preferences and measures information about the environment in these characteristically *internal* states. The information transfer model in Chapter 7 is technically rather difficult, I try to explain its implications in simpler terms at the end.

Part III: To reiterate, even a purely descriptive account of normative intuitions and language can help constrain cultural relativism with respect to them. Chapter 8 develops a model of primitive meaning content, inspired by the teleosemantic theory of Ruth Millikan (1984) and a game theoretic model from Brian Skyrms (1996). This general model is then applied to regulatory hierarchies, resulting in systems with multiple semantic maps sharing many formerly puzzling features of human normative deliberation. Chapter 9 defends the plausibility of this primitive content hypothesis against standard objections. In particular, the "open question argument" of G.E. Moore and Hume's famous analysis of normative relationships do not apply to historical-functional semantic theories are addressed, along with a bakers dozen (or so) other philosophical worries. The reason why you can't get an ought from an is is explained.

Throughout, I try to keep focused on the theoretical objectives involved in an evolutionary theory of knowledge. This means resisting the temptation to delve into "hot" topics like the nature of consciousness, group selection, the excesses of adaptationism, the marginalization of developmental biology, and the naturalization of ethics. I apologize in advance to readers who find this stinginess unsatisfying.

Pre-established Harmony

German philosopher Gottfried Wilhelm Leibniz (whose most famous accomplishment may well be the invention of the differential calculus simultaneously with Newton in the closing years of the 17th century) argued that substances exist but do not interact, each separate thing sufficient in itself to determine its unfolding over time in the absence of

causal interactions with other things. One consequence of this view was that our knowledge of things in the world cannot be the result of causal interactions with them, as we usually assume. Leibniz' solution to this epistemological difficulty was that God, in creating each thing ("monad"), had done so in such a way that the unfolding of each thing over time was mirrored by others. Thus, perception and knowledge are not the result of causal interaction, but of a "pre-established harmony" between the movements of our minds and the movements of things in the world.

Scottish philosopher David Hume, like Leibniz, called into question our ordinary notion of causality, but as an empiricist, both his method and his conclusions were different. Hume asked us to consider closely, not conceptual analysis, but the act of perceiving causal interactions. He simply pointed out that all we ever see are sequences of events, never the actual force of one object acting on another. Causal power is something we project onto interactions, rather than perceive in them, and the notion of causality itself, he speculated, was merely the generalization of our own expectant impulses. Despite, or perhaps because of, Hume's informality and reliance on common sense, his argument remains to this day *the* problem of causality, a large part of his not inconsiderable philosophical legacy. Along with the "problem of induction," his argument concerning causality cemented his reputation as one of the greatest skeptics and philosophical troublemakers of all time. What is usually disregarded is his own proffered solution to the difficulties he raised.

Hume argued that causal and inductive reasoning, along with moral judgement, are not the result of the rational perception of eternal laws of reason and standards of behavior, but are in each case the result of instinct, habit, or sentiment. In contemporary terms, Hume believed that we are just "wired" up to think the way we do, and that rationalist notions of being able to perceive certain relationships as self-evident by the "light of nature" are just nonsense. This was not intended to undermine our reliance on causal or inductive reasoning nor our genuine belief in moral standards, however, but to pull the rug out from under rationalist pretensions that Reason was the final authority on all aspects of human thought and behavior. But how, without recourse to God or the power of Reason, were we to establish the propriety of causal and inductive inference?

Hume did know when to give it a rest. Philosophy, he said, takes one into distressingly deep waters and forces one to conclusions that seem to undermine everything one believes. But fortunately, human nature is too strong for mere philosophical arguments to freeze us into inaction, and one needs to know when it is time to put down the philosophy and go play billiards. This pragmatic streak shows itself in this "solution" to his skeptical doubts, as well. It may be that causal and inductive inference are merely the operation of instincts or habits. Nonetheless, any fool can see that they are *good* instincts and habits, without which we would be quite incapable of surviving. Causal and inductive reasoning seem to "fit" with the patterns of the world in just the right way to help us cope with them. How this could be so was not clear, but that it is so only a philosopher could doubt. What he said, was this:

> *Here, then, is a kind of pre-established harmony between the course of nature and the succession of our ideas; ... As nature has taught us the use of our limbs, without giving us the knowledge of the muscles and nerves, by which they are actuated; so has she implanted in us an instinct, which carries forward the thought in a correspondent course to that which she has established among external objects; though we are ignorant of those powers and forces, on which this regular course and succession of objects totally depends.* ***An Enquiry Concerning Human Understanding****,* § V.

The reason for bringing up Hume's little joke on Leibniz is that Hume was right. Nature *has* implanted in us cognitive instincts which keep our thoughts in productive harmony with the world, and that is the secret to understanding knowledge. What we know and Hume could not is how in fact nature has gone about establishing this harmony between cognition and worldly processes. One hundred years before Darwin, Hume's confidence in our reasoning processes was merely common sense, without theoretical foundation. For us, over two centuries later, a great deal is known both about the broad outlines of evolutionary theory and about the details of our own evolutionary history. Much work remains to be done in both areas, and the empirical task of working out the details of our evolutionary history may well never be completed, due both to the enormity of the task and the antiquity of the relevant facts. Theory can, at least temporarily,

provide a little more closure.

This book, then, is a contribution to the theory behind Hume's pre-established harmony. Insofar as it is successful, it will do what successful philosophy often does — take its subject matter a little closer to being an autonomous science with its own distinctive methods, a little farther from being the proper domain of philosophers.

William Harms, Seattle, 2001