Gregory Nirshberg
University of Texas at El Paso
gnirshberg@miners.utep.edu

**Science, Progress and Evolutionary Epistemology**

*"The most puzzling feature of science is that it works so well in realizing its manifest goals, so much better than any other social institution"* – David Hull (Campbell, 1988)

## Introduction

The endeavor of science has turned out to be humanity's most successful way of coming to understand the universe we live in. Science and scientific theories have allowed us to build computers, planes, and bridges. We can eradicate diseases and remove and replace organs. We've sent man made artifacts out beyond the confines of our planet. But how is it that science works so well? And can we apply the principles that allow science to work so well, to a general theory of knowledge? Both in science and in epistemology we run into the skeptical problem. How is that any of our beliefs can correspond to some state of affairs in the world, if we have no direct access to that state of affairs? What external observer can confirm that our beliefs are "True"?

In the following pages, a defense of evolutionary epistemology as a general theory of knowledge will be undertaken. Using various examples from the history of science, I will discuss what it is about science that allows it to work so well, and how that corresponds to the principles that underlie an evolutionary epistemological framework. This will require a shift towards a more pragmatic conception of both knowledge and science (not surprising as the pragmatists were heavily influenced by evolutionary theory). It will also be both a descriptive endeavor to explain knowledge acquisition, and a prescriptive endeavor of exploring how we can improve this process through the insights possible when accepting a version of evolutionary epistemology.

## Naturalized Epistemology

There are a variety of epistemological theories that fall under the label of "naturalized" epistemology. What all these disparate theories have in common is a recourse to the sciences to explain, or help explain, what justification we have for our knowledge. Weak versions of naturalized epistemology tend to make the relatively uncontentious claim that the information coming out of the sciences which is relevant to human cognition and reasoning, and that describe how humans come to form beliefs about the world, should be integrated into our epistemological pursuits and attempts to resolve epistemological questions. I claim this is uncontentious in the sense that Philip Kitcher means it when he says, "How could our psychological and biological capacities and limitations *fail* to be relevant to the study of human knowledge?" (Feldman, 2001) The problem with this approach is that it's not a theory in and of itself, but simply a comment on the fact that whatever preferred epistemological theory we do have, should take into account data and research coming out of the cognitive sciences that is relevant.

A more extreme version of naturalism, known as replacement naturalism, was heavily influenced by W. V. Quine. This stronger version advocates entirely replacing the pursuit of epistemology with a scientific study of how we reason and come to form beliefs (Quine, 1969). Quine wrote of replacing epistemology with psychology, though today we would likely speak of replacing epistemology with (a branch of) cognitive science. Quine thought there was nothing more we could do with epistemology than the descriptive endeavor that psychology allowed us to do. There are two normative problems that stem from this prescription though. The first is that in advocating the replacement of the normative pursuit of epistemological questions, with a descriptive endeavor, Quine is implicitly making a normative statement. But if Quine turns his own advice back on his prescription, then he would have to admit that his own statement lacks any normative support, and circularity ensues. Another problem is that in saying that the cognitive sciences can tell us anything useful about human reasoning and belief formation is already accepting the normativity of these sciences. But if we can't make any normative epistemological claims about knowledge, then what justification do we have for thinking anything that the cognitive sciences (or any sciences) tell us is true in any way? This is precisely the issue we want to understand.

Besides these normative issues, the root problems and criticisms of only being descriptive remain. Evolutionary epistemology, with tends to fall under the broad umbrella of naturalized epistemologies, thus is hit with this criticism as well. And the spirit of the criticism, if not the truth value of it, is an important one. If we are concerned with addressing epistemological questions of what counts as knowledge and how we can have it, our theories need to have normative import. As I pointed out in the previous paragraph, normativity is unavoidable, even when explicitly trying to avoid it. So in the arguments that follow, evolutionary epistemology will be presented as both a descriptive theory falling within the naturalized epistemological framework describing how human reasoning and belief formation result from certain types of natural processes, as well as a prescriptive theory of how we can best go about forming true beliefs about the world we find ourselves in.

**Evolutionary Epistemology**

Evolutionary epistemology consists of two broad programs. The first has been dubbed the EEM program, or the evolution of epistemological mechanisms program. The second has been dubbed the EET program, or the evolutionary epistemology of theories program (Bradie & Harms, 2008). The former consists of the claim that biological evolution by natural selection has been the generator of reliable perceptual and cognitive mechanisms which result in a certain "fit" between those mechanisms and the world. This thus gives us justification for saying that our perception of the world is relatively accurate, as well as our other propositional beliefs about the world, because if this weren't the case, these structures would have never made our survival throughout evolutionary history possible.

The EET program abstracts away from the specifics of biological evolution, using analogies and metaphors as a way to describe the evolution of our ideas and theories about the world. Many working within an EET framework will say that at the root of all knowledge formation is a general process of variation and selection, and that this process can gives us justification for everything from holding certain beliefs to working with certain scientific theories. While some who advocate for an EET program simply use evolution as an analogical framework to discuss epistemological issues, others take the stronger stance that evolution by natural selection is itself actually a special case of a more general kind of process in the universe (Bradie, 1986), and biological evolution just happens to be the one we are most familiar with and have the most knowledge about. When we abstract away from the specific biological mechanisms we are actually describing this more fundamental process. What this entails is that EET theories are not required to make explicit analogy to every individual aspect of biological evolution for the EET program to be successful. For instance, evolutionary epistemology doesn't necessarily have to account for genetic drift or random mutation (criticisms that are too often leveled at EET) to provide a valuable descriptive and prescriptive mechanism for belief formation.

As originally formulated by Donald Campbell, what is integral for both biological evolution and for belief formation or problem solving is a mechanism for introducing variation, some sort of selection process, and a retention mechanism. Campbell's evolutionary epistemology was thus labeled a VSR model, a variation/selection/retention model of epistemology (Christensen & Hooker, 1999). Explicating precisely what is varied, what is selected for, and what the retention process has been a difficult task though. Biological evolution has the advantage of a very large field devoted to explicating these issues, and while the specific unit of selection is still debated, in general, genetic mutations introduce variation in the population, selection consists of survival of the organism in service of reproduction by the organism, and what is retained is the genetic code that builds the eventual new organism phenotype.

EET models must abstract away from biological evolution and in laying out this more general process, must discuss precisely what it is that is being varied, what is being selected, and how it is retained. The model I will be laying out will explicate these details both in reference to biological evolution, and in reference to knowledge acquisition. The following attributes of an evolutionary epistemology model of knowledge will be discussed: 1. Action selection and the satisfaction conditions inherent in it. 2. Interactive opportunities and the possibility of error. 3. Openness to change and error correction. 4. Evolution *and* knowledge acquisition as a process of moving away from error. The connections between all of these points will be made clear in the following pages.

**Action Selection**

Action selection seems like a difficult notion to unpack in regards to non mobile biological system such as plants, but this actually isn't the case. While plants don't engage in behavior in the way we associate with organisms that can move, plants do have a particular phenotype which results in a system with a certain interactive relationship with its environment. Thus in the case of a plant, the entire constructed phenotype and resulting relationship with its environment can be viewed as one rigid form of action selection. There are satisfaction conditions in this rudimentary phenotypic action selection, in that the constructed phenotype is in a sense *implicitly presupposing* it is in a particular environment where its phenotype will successfully reproduce. If the environment changes, or there was a genetic mutation resulting in a phenotype not suited for the particular niche it finds itself in, the phenotype will fail to achieve those satisfaction conditions, by being less likely to reproduce.

Once organisms can move and interact in their environments in a more active way, action selection becomes more important and we can talk about the satisfaction conditions of that action selection in a much more explicit way. A bacterium that is swimming up a nutrient rich sugar gradient has selected an action with the satisfaction condition of providing energy to the system. It can achieve or fail to achieve that satisfaction condition, by being in an environment that does or does not contain sugar. For instance, these bacterium will swim up sugar gradients and when they find themselves not swimming up a sugar gradient will initiate tumbling and then stop and proceed to swim again (Bickhard, 2008). If the bacterium is swimming up a saccharine gradient though, it cannot engage in fine grained enough discrimination to determine this fact, and engages in swimming action selection, thus failing to achieve the satisfaction condition of acquiring its needed energy.

While survival or not surviving is the end result of a lifetime of action selections, action selection itself is about the instrumental needs of the organism in the moment of action selection. This is true of simple biological organisms, and it is true of human beings. Action selection stems from the needs, desires, and beliefs of the organism interacting with the environment.

**Interactive Opportunities**

Any agent is limited by its neurophysiological structure and the environment it finds itself in as to the kinds of interactive opportunities available to it. Whether a plant or a bug, or a cat or a human, to what degree we are mobile and the kinds of sensory systems we have, determine the kinds of possible epistemic access we have to the world. Our epistemic access to the world creates a bit of a paradox for biological systems. The broader our epistemic access is to the world, based on our interactive opportunities, the broader are our possibilities for having false beliefs about this world. If your world consists of sugar and not sugar, the only times you can be said to have a false belief is when you encounter saccharine. It doesn't matter that there are a range of objects that fall under the title of 'not sugar' that could possibly be discriminated as individual (had you a different set of sensory systems); for your needs, it simply isn't relevant. The broader the range of interactive opportunity, the greater the ability to make fine grained discriminations about your environment, and the more chances there are for being in error based on your satisfaction conditions. This is a valuable thing though, because the more possibilities there are for error, the more chances we have to detect that error.

This is implicitly accepted within science, and a quick example will help make the case. Imagine a clinical drug trial that is done with 100 white males between the ages of 25-35, and a certain set of results are obtained about the efficacy of the drug, possible side effects, etc…let us say the drug is shown to be efficacious and there are no side effects. This would normally be taken as relatively good evidence that the hypothesis that the drug is efficacious and has no side effects is valid. If the same experiment is done again with another 100 white males of the same age group, and the same results are obtained, this would likely lend more credence to the hypotheses, but only marginally so. And every time the same experiment is performed with the same conditions, less and less support for the hypothesis will be gained (Achinstein, 1998). But if the experiment over time is successively widened to include 1000 people, and then people of both sexes, then all age groups, and eventually all races, if the results continued to obtain, this *would* lend more significant support to the hypothesis. This is not news to anybody. But what is relevant is the way in which we interpret why this is so. The argument presented here is that this is opening up the hypothesis (belief), which is tested by experiment (action), to broader and broader interactive opportunities (experimental conditions), allowing for more possible epistemic access, opening itself up to more possibilities of error. As a general process, this is just a first step. To correct for error, the error must be known first, and so then there needs to be a mechanism to account for this error correction.

**Openness to Change**

At the root of all possible variation in evolution is an openness to change (Christensen & Hooker, 1999). Again, this sounds odd to speak of for simple biological systems, so let's flesh it out. In simple biological systems there is no change throughout the life of the organism. A plant will have a certain phenotype, and besides its changing morphology (which admittedly is a type of change that will have an effect on its ability to survive and reproduce), it cannot change in ways that allow it to better select actions to engage in its environment with (again, besides changing morphology). For these types of biological systems, openness to change simply refers to genetic mutation in the reproductive process. The organism itself does not change, but the overall temporal process of reproduction over the course of generations results in different phenotypes because of the fact of genetic mutation. As organisms evolve to the point where synaptic plasticity occurs, openness to change takes on a broader meaning, not only do phenotypes change over successive generations, but the behavior of an organism can change throughout the course of its life based on its changing neurophysiology, allowing for skill acquisition and learning of all sorts (learning to seek certain things, learning to avoid certain things, etc…). The ability to learn how to learn, to subject the learning process to its own process of variation and selection is also another important evolutionary development (Christensen & Hooker, 1999). We humans can take these to an entirely new conceptual level of being able to understand that we can be in error, and that there are ways in which we can change our behavior to account for this error, to correct for it. But only when you become aware of error can the possibility of error correction occur. Knowledge of error always precedes knowledge of how to account for error (Bickhard, 2002).

What we end up with is the following general process. Organisms engage in action selection, with certain satisfaction conditions built into that action selection, implicitly or explicitly. Thus organisms can be in error based on their needs. The possible interactive opportunities that an agent has can be looked at as its possible epistemic access to the world, and from this stems its possibility for error detection. All agents have some level of being open to change (over the generations or throughout their lives) and this openness to change allows for error correction. And so what we have is a model of how organisms engage in a process of moving away from error. For simple biological systems this is an evolutionary story, for us, I am arguing that this is a process described by a variation/selection/retention model of evolutionary epistemology.

This notion of moving away from error addresses multiple classically problematic aspects of science, two of which end up having an interesting connection. The first is negative induction, the idea that our scientific theories are fallible, and that each and every theory in the past has proved to be wrong. If this is the case, what warrant do we have for thinking our current theories are correct, or that any theories in the future will ever be correct? The second has to do with Thomas Kuhn's assertion that scientific theories progress through revolutions and paradigm shifts (Kuhn, 1999). If scientific theories aren't built on cumulatively, but are completely wiped out and replaced with different theories, different paradigms, how can scientific knowledge be said to be progressive? So this ends up being a problem both for warrant for beliefs, but also for scientific realism. Looking at these changes through the lens of a process of moving away from error, we see that when a scientific theory replaces an older one, it does so because it is in error under fewer (known) conditions than the previous theory. The old theory couldn't account for a piece of evidence, or failed to make a proper prediction. And the new theory is able to account for an anomaly, make better predictions, provides a better explanation, allows for more practical applications and control over the environment, etc…

Our scientific knowledge, positive knowledge in general, IS progressive, but not necessarily cumulative (Bickhard 2002). It's progressive only in the sense that it is able to account for more error. Moving away from error is progressive *and* cumulative though. And this is why whether scientific revolutions really happen, and whether science progresses through paradigm shifts, isn't actually what is important to focus on. Either one happens, or it doesn't. In either case, our knowledge wasn't actually revolutionized, but, rather, we have been able to cumulatively account of more error. Sometimes in accounting for more error a change in paradigm is necessary, other times it isn't, but science doesn't have to progress by paradigm change as long as error knowledge has increased, and thus the argument about scientific revolutions becomes largely moot.

**Criticisms and Responses**

There are a handful of criticisms leveled against EET versions of evolutionary epistemology that are worth considering. These criticisms are worth considering for a few reasons. 1) When viewing evolutionary biology in the proper way, the criticisms actually point not to the weaknesses in evolutionary epistemology, but rather, its strengths, and 2) When fleshed out, the parallels between these aspects of evolutionary biology and the field of science (and progress within it) become even more clear.

**"Evolution isn't about optimal fit, thus if you use evolutionary epistemology as a way to justify knowledge, you can't explain the correspondence between our beliefs and the state of affairs in the world"**

Natural selection does not produce organisms that have some sort of optimal fit or correspondence between their knowledge and some sort of static pregiven state of affairs in the world, correct. Natural selection does not select what is most optimally suited to the environment. Rather, natural selection discards what is least suited to the environment; it doesn't select the best, but prunes the worst. Natural selection guarantees only that what survives is viable (Varela, Thompson, Rosch, 1991).

Implicit in this is the notion, under the framework that has been developed here, that the organisms that don't make it, that don't successfully reproduce, behave in ways that are in error under more conditions relative to the constraints of their environment than other organisms. Organisms that survive will tend to behave in ways that avoid the errors of the ones that don't behave in those ways.

And when mutation occurs that leads to phenotypes that vary in their behavioral tendencies, the ones that tend to behave in ways that avoid more errors will, again, be the ones that tend to reproduce. Evolution is fundamentally a process of selecting out organisms that fail to avoid error, allowing the mutations that occur to construct organisms who are better and better at avoiding error in the long run (remembering that error is being defined as failing to achieve the satisfaction conditions of action selection, whether at the level of behavior or of the phenotype in general).

Biological systems can thus be viewed as knowledge systems. Their behavior as theories that can succeed or fail, and the process of evolution as a process of gaining knowledge (Bradie, 1986). We have to remember that knowers and biological systems are not passive recipients of data from the environments, but active agents in those environments. Charles Sanders Peirce talked about knowers as agents, "who obtain empirical support for their beliefs by making experimental interventions in their surroundings and learning from the experiences that their actions elicit." (Hookway, 2008) Belief formation is not only the result of interaction with the environment, but beliefs (behaviors and the anticipations inherent in them) go on to have further causal interactions with the environment, thus providing further input for belief formation.

The problems of correspondence between beliefs and the state of affairs in the world is a problem for all epistemological theories, not just evolutionary epistemology. We can never be justified in asserting an absolute correspondence between our beliefs or scientific theories and the state of affairs in the world, as there is no omnipotent philosopher or scientist in the sky to confirm our beliefs. The fact that evolution isn't about optimal fit, but rather about selecting out what is not viable, allowing for a process of fallible error correction, is actually what makes biological evolution so powerful, and what makes evolutionary epistemology so successful as well. Again, Peirce had the correct notion when he argued that the focus of epistemological inquiry shouldn't be in showing how we could possess absolute certainty, but in understanding “how we can possess methods of inquiry that contribute to our making fallible progress.” (Hookway, 2008) Whether our beliefs correspond to reality is never something we can say with absolutely certainty, but evolutionary epistemology gives us warrant accepting certain things to be true, while continuously engaging in the process of fallible progress.

**“Biological evolution is blind, it is random, and thus it serves as a flawed analogy for knowledge construction”**

Another criticism that is sometimes leveled is that evolution is blind, a random process, and thus the epistemological analogy is flawed. This is quite a confusing criticism, since it seems that it is precisely the fact that EET is NOT a blind process that makes it so valuable and allows us to transcend the limitations of a blind variation and selection process. The fact that the acquisition of scientific knowledge is directed rather than random, on its own, says nothing significant about the analogy, as we must remember that we are abstracting away from the specifics, to describe a more general process.

The criticism is also somewhat flawed in another way though. Biological evolution is not really completely random and blind, rather, it is a process that is open to novelty, as defined earlier. Selection only takes place among the actual phenotypes as a result of genetic mutations that appear in the environment. And these variations are variations on the current acquired behavioral patterns (beliefs) about the world. Thus selection is always built off previous knowledge, but it is varied in a trial and error fashion, rather than directed (Bradie, 1986). Directed variation does not contradict the variation and selection paradigm, but is what makes human knowledge formation so strong. While too much trust in the presuppositions of the current state of knowledge can lead to problems at times, in general, being “theory laden” and guided is a central expression of science’s cognitive power, not an epistemic defect (Christensen & Hooker, 1999).

What's powerful about this is that the variation and selection paradigm applies to the very process of variation and selection. Thus in science, not only do theories go through a process of variation and selection, but the very criteria for theory choice do as well. Popper saw this when he made the point that critical discussion about theory choice can only stem from the best tested theories at any given time. And to defend why we should put any faith in the idea that this critical discussion justifies theory choice, argued "We adopt critical methods which themselves have withstood severe criticism." (Salmon, 1998) Popper's point was that the very values and considerations that go into choosing one theory over another also change over time as a result of empirical matters. Kuhn had proposed five unchanging values that are used in theory choice: accuracy, consistency, scope, simplicity, and fruitfulness (Kuhn 1998b). Ernan McMullin (1998) responded to Kuhn and asked why? Why must these shared values remain an unchanging permanent feature of science? If other criteria for theory choice change, why not these? He further argues that we can justify our epistemic values by appeals to history and experience (McMullin, 1998). As an example he discusses the change in our valuing simplicity and fertility after the Copernican revolution; that it was through experience that these values were found to be empirically useful. Whether scientists embraced these values knowingly in the moment, or only as the result of an analysis with the benefit of hindsight, is not necessarily relevant. An openness to novelty, even accidentally, can allow for progress (as it does in biological evolution).

We must be careful though, while some values may be acquired through empirical means, not all are. And even the ones that are, might not be perfect (similar to how better scientific theories that replace older ones are themselves replaced). We must consider that we may have a certain value that we hold simply because it stems from biological factors that arose through evolution. We have what Helen de Cruz and Johan De Smedt dubbed intuitive ontologies (De Cruz & De Smedt, 2006). These ontologies, if they are the result of natural selection, only need to be adaptive, and so the knowledge they provide is not guaranteed to be a genuine reflection of a causal mechanism in the universe.

Whether it be these intuitive ontologies or other cognitive biases, there are many ways for our methods of inquiry to lead us astray. Given the nature of our particular evolutionary history, Paul Churchland pointed out that it would be “miraculous if human reason were completely free of false strategies and fundamental cognitive limitations, and doubly miraculous if the theories we accept failed to reflect those defects.” (Thomson, 1995) Our understanding of things like confirmation bias, our tendency to give our own theories more weight, and the feminist critiques against science are all important reminders that science is an activity engaged in by people, fallible human beings, and that the results of the process engaged in by these individuals are influenced by whatever cognitive faculties the individual brings to the table. This is yet another reason why ‘openness to novelty’ is important. By embracing falliblism, and the idea that we can improve, we can be on the lookout for ways in which our values and biases lead us astray. Only when the scientific community, and individuals within it, become aware of things like confirmation bias, or that certain unjustified intuitive ontologies have been influencing science, or that women have been marginalized, can they be accounted for. The same openness to novelty that makes evolution so successful (by way of genetic mutation on through to synaptic plasticity), extended to a cognitive reflective process, is what makes science so successful as well.

**“Environments are always changing, there is no progress towards a goal in evolution, just adaptation for different niches, but in science there is a goal, thus the analogy is flawed”**

A strong criticism often leveled surrounds the nature of organism/environment interaction in evolution. Accepting that evolution is not about optimal fitness, but about viability, and further accepting that environmental conditions are always changing, an important point follows. Evolution results in organisms constructed to survive in particular niches present in the environment. There is no progress towards a specific goal in evolution, only viability in the environment. As environments change, organisms change to fill the new niches afforded in that environment, but no more. The criticism that follows is that science on the other hand, IS progressing towards a goal. So not only is the analogy flawed, but is fundamentally describing something at odds with the goals of science.

There are two ways to respond this criticism. Both of these ways admit the factual nature of the criticism, while denying it poses a problem for evolutionary epistemology, by questioning what the goal of science can legitimately aim to be. Biological evolution happens to active agents, but because the process itself does not have a goal, what is selected for is instrumental for survival value in the niches afforded by the environment. But most human beliefs, and especially scientific theories, are not instrumental in a survival sense. How do we flesh out the relationship between human beings and scientific theories in relation to a goal? Simply put, in the same way as all other organisms and the environment. Individual scientists and even the scientific community may have a goal of “truth”, but we’ve already seen both in evolution, and in the history of epistemology, there is this problem of the correspondence relationship between truth and our beliefs about the world. If that’s the case, then we can’t know beforehand what the goal of science is, besides in the abstract. We don’t have a blueprint describing what the truth is, in which case we can’t have any infallible knowledge of how to reach it, or even *when* we’ve reached it. We’re never in a position to recognize whether any of our theories are “True”, all we can say is that it meets the standards of acceptance that are endorsed, for the time being, in the scientific community (Hookway, 2008). It fills an instrumental niche in the community of practitioners based on their needs and interests.

Just as organisms and their environments shape and reshape each other, so it is the same with humans and scientific theories and beliefs. What is the instrumental value of science? Science allows us to control and manipulate our environment through the use of engineering and technology. It allows us to predict the future and explain the past, and all of that is in service to action selection. As John Dewey noted, “In some sense, all inquiry is practical, concerned with transforming and evaluating the features of the situations in which we find ourselves“ (Hookway, 2008). And again, all action selection has satisfaction conditions and falsification criteria. But the question follows, why should we assume that control and manipulation of the environment, and the scientific theories that these practical uses stem from, are progressing towards any limit of truth; are progress in any way? It’s worth stepping back for a second, but before doing that, let’s draw attention to an important point. Even most fallible accounts of knowledge will admit that there is an objective world out there, and the issue is our epistemic access to it. Similarly, in biological evolution, while the environments an organism finds itself in may change, the laws of physics that the environments are constructed as a result of don’t change.

**Scientific Realism or Instrumentality?**

If we take the intentional stance towards an organism, when an organism choose a particular behavior to engage in, it’s presupposing that it’s in a certain environment where that behavior is appropriate based on the anticipated satisfaction conditions of engaging in that behavior. Only if the selected behavior fails to achieve satisfaction conditions can an organism have a notion that the belief or behavior was in error. And only when you are aware of error can you account for it in knowledge. In the moment, before that awareness occurs, something can be held as true, and only through further interaction can it be shown to be false. This is especially true for something like science where there is no external observer, no omnipotent philosopher who can point to whether the belief corresponds to reality at all. But that doesn’t mean we can’t have warrant for holding a belief as a temporary or conditional truth in the meantime, otherwise we would be unable to act in the world. In fact, the *only* way to *know* whether things are true is to interact as if they are, and see whether our interaction is what we expect or not (try to manipulate or control something and see if it works). Organisms have to act in their environments to survive, and we need to behave as if certain beliefs we hold are true.

The same holds for science. Hypotheses are akin to provisional beliefs. We act *as if* the belief is true by engaging in the action selection that will fulfill the anticipation conditions of that belief; in the case of science this involves experimentation. And those experiments can achieve or fail to achieve the satisfaction conditions of that action selection (experiment). General relativity provides a useful example. When proposed, general relativity accounted for the anomalous perihelion of mercury, as well as other things, but lacked a strong empirical foundation at first (Curd & Cover, 1998). But it was able to make certain predictions. And the way I want to view those predictions is that, in a sense, behaving as if the theory was true contained certain anticipations in it. And there were certain satisfaction conditions for those anticipations, and certain ways those satisfaction conditions could fail to hold. Experiments can be viewed as engaging in action selection, and in the case of the theory of relativity, the satisfaction conditions held, and we were able to provisionally accept general relativity as true, even though we know it can't be the whole story; that it will eventually be overturned. When a theory survives falsification, it just tells us it was good enough, but not perfect. As Popper pointed out, we can never prove a scientific theory true, we can only fail to falsify it. (Popper, 1998) The behavior engaged in, or the theory we have about the world, may not be fine grained enough to turn up errors (and in general is underdetermined by the data).

The more possibilities for interaction that there are though, the more possibilities for error there is. The more possibilities for error, the more possibilities there are for error detection and error avoidance and correction. The more errors we correct the more knowledge can be said to progress. Our knowledge is necessarily wrapped up with our epistemic access to the world and our possibility of being able to uncover that we are wrong. If Newtonian mechanics let us interact in the world without any errors in prediction (whether in a failure to predict something, or in a false prediction), if it allowed us to create technology based off its laws that worked perfectly (besides human error), then we would have every reason to assume it was true, even if an outside observer could see that we were wrong. We have to have some sort of epistemic access to the fact that there is something more out there, something has to go wrong with the theory, otherwise there is no need to vary it, and thus nothing to select. This is both a theoretical argument and a practical argument. Interaction is not just the kinds of basic interaction we think of when describing simple organisms, I've already mentioned technology and engineering as examples. Our scientific theories allow us to create artifacts that allow us to interact with the world and with them. Bridges are built on certain principles. Those bridges either remain structurally sound or don't. If they don't, something about the principles under which we constructed them around are in error. But importantly, even if they do remain sound, all that tells us is that we can provisionally accept it as knowledge. The range of interactions we've subjected it to might not have been tested all the possible interactions, or our error falls within a range of acceptability. And only through ongoing interaction can the truth of this be determined.

The tools we create allow us to interact in new and different ways, which again open us up to the possibility of error. Our artifacts determine and limit the range of tests we're physically capable of putting something to, or the conditions under which we can test a hypothesis under. It is these tools, and our imaginations that allow us to transcend the limitations of the niches most organisms fill in their environments; we in a sense create our own niches, the environment doesn't just push back at us, we push back at the environment. Though we are limited to what is currently possible at any given time given our current base of knowledge, and so our ability to broaden our interaction possibilities is itself subject to the constraints of the evolutionary epistemological process.

**Conclusion**

Traditionally in epistemology there has been a division between how we come to know true things about the world, and what it is that makes something true. Evolutionary epistemology (and process approaches in general) remove that distinction. Our lack of epistemic access to the truth forces the fact that knowledge is a constructive process, and one that changes over time. Thus, how we come to know true things about the world IS what makes them true, it is the process that is fundamental, not the proposition. We can never be certain about the correspondence between our beliefs and the state of affairs in the world, but we *can* make fallible progress by accounting for more and more error in our interactions with the world.

Our current state of knowledge determines action selection and their satisfaction conditions based on the relationship between our knowledge and its instrumental role in our relationship with the environment. The possibility of error always exists in any interaction. But what allows error correction, what allows progress is generally speaking, an openness to novelty and a broadness of interaction possibilities.

In science, our anticipations reach their satisfaction conditions more and more regularly, which tells us that the process that we have been implementing to construct those anticipations has positive normative value. If this is the case, then we can abstract away from this the mechanisms that are at work that have allowed this endeavor to be successful. By developing a general theory that adequately describes this process, we can apply the principles of that theory back to the process itself, as a normative proscriptive endeavor of how to improve the process which we have to some degree been engaging in blindly without an awareness of the principles that were making it work. In this way we can not only say what it is about science that has worked so well, we can apply those very principles back to the scientific endeavor itself. We can remove aspects of the process that don't confirm to these principles, and we can incorporate more that do. And if the result of making these changes is somehow a loss in the efficacy of science, then we can further critically examine the changes we made and determine how and why they went astray. In this way, we make fallible progress by engaging in this process of variation, selection, and retention.

# References

Achinstein, P. (1998). Induction, Prediction, and Evidence: Explanation v. Prediction: Which Carries More Weight? In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 481-493). New York: W. W. Norton & Company

Bickhard, M. H. (2002). Critical Principles: On the Negative Side of Rationality. *New Ideas in Psychology, 20,* 1-34.

Bickhard, M. H. (2009). The interactivist model. *Synthese*, *166*(3), 547-591.

Bradie, Michael (1986). Assessing evolutionary epistemology. *Biology and Philosophy.* 1 (4):401-459.

Bradie, Michael and Harms, William, "Evolutionary Epistemology", *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2008/entries/epistemology-evolutionary/>.

Campbell, D. T. (1988). A General 'Selection Theory', as Implemented in Biological Evolution and in Social Belief-Transmission-with-Modification in Science. *Biology and Philosophy.* 3:171-177.

Christensen, W. D. & Hooker, C. A. (1999). The Organization of Knowledge: Beyond Campbell's Evolutionary Epsitemology. *Philosophy of Science.* 66: 237-249.

Curd, M. & Cover, J. A. (Eds.). *Philosophy of Science: The Central Issues.* New York: W. W. Norton & Company

de Cruz, H. & de Smedt, J. (2007). The role of intuitive ontologies in scientific understanding – the case of human evolution. *Biology and Philosophy.* 22 (3):351-368.

Feldman, Richard, "Naturalized Epistemology", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/epistemology-naturalized/>.

Hookway, Christopher, "Pragmatism", *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2010/entries/pragmatism/>.

Kuhn, T. (1998a). Rationality, Objectivity, and Values in Science: The Nature and Necessity of Scientific Revolutions. In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 86-101). New York: W. W. Norton & Company

Kuhn, T. (1998b). Rationality, Objectivity, and Values in Science: Objectivity, Value Judgment, and Theory Choice). In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 102-118). New York: W. W. Norton & Company

McMullin, E. (1998). Rationality, Objectivity, and Values in Science: Rationality and Paradigm Change in Science. In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 119-138). New York: W. W. Norton & Company

Popper, K. (1998). Science and Pseudoscience: Science: Conjectures and Refutations. In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 3-10). New York: W. W. Norton & Company

Salmon, W. (1998). Induction, Prediction, and Evidence: Rational Prediction. In Curd, M. & Cover, J. A. (Eds.), *Philosophy of Science: The Central Issues.* (pp. 433-444). New York: W. W. Norton & Company

Thomson, P. (1995). Evolutionary Epistemology and Scientific Realism. *Journal of Social and Evolutionary Systems.* 18(2): 165-191.

Varela, F., Thompson, E. & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge: MIT Press.