

Toward a Pragmatic Account of Scientific Knowledge

Jeffrey A. Barrett
7 October 2001

Abstract: C. S. Peirce's psychological analysis of belief, doubt, and inquiry provides insights into the nature of scientific knowledge. These in turn can be used to construct an account of scientific knowledge where the notions of belief, truth, rational justification, and inquiry are determined by the relationships that must hold between these notions. I will describe this account of scientific knowledge and some of the problems it faces. I will also describe the close relationship between pragmatic and naturalized accounts of scientific knowledge.

I. Introduction

On the standard account of propositional knowledge an inquirer K knows proposition P if and only if (i) K believes P, (ii) P is true, and (iii) K has rational justification for believing P. Given this, one might try to construct a theory of knowledge in three parts: explain what it is to believe a proposition, explain what it is for a proposition to be true, and explain the conditions under which a belief is justified, then take one's theory of knowledge to be the three component theories considered together. But if the three component theories are constructed independently, there is no reason to suppose that they will mesh to form a single, coherent account of knowledge; rather, one has every reason to suppose that each of the three component theories will properly constrain what one might consistently say in the other two theories.

One relation that constrains the notions of truth and rational justification is that rational justification should track truth. Rational justification might track truth perfectly by guaranteeing the truth of a rationally justified belief (as Descartes claimed that clear and distinct perception guarantees truth), or rational justification might track truth imperfectly by making it more likely than not for a rationally justified belief to be true (as reading in the morning paper that it will rain justifies the belief that it will rain). In either case, if rational justification tracks truth, then this explains why having rational justification for a belief provides good reasons for holding the belief to be true (or probably true or probably approximately true). And this explanation is something we presumably want from a satisfactory account of truth and rational justification since if one does not understand why one's criteria for having rational justification should track truth, then it is not at all clear why knowledge should involve rational justification at all.

This relationship between truth and rational justification might also constrain one's account of belief. If having rational justification for a belief does not guarantee its truth (as modesty, good sense, and scientific practice suggest), then having rational justification fails to justify full belief. If justification tracks truth only to a degree, then justification is presumably only justification for a degree of belief. So if justification falls short of guaranteeing truth, then perhaps one can only capture the spirit of the truth connection by adopting an account of belief that allows for degrees of belief.

If scientific justification and belief typically come in degrees, then an account of knowledge where an inquirer either knows or fails to know each proposition is clearly inadequate to the task of even describing scientific knowledge. Rather, we presumably want an account of scientific knowledge where inquirers believe to a degree that is

proportional to whatever degree they believe that they have good reasons for believing. But how belief is in fact established is a subtle matter.

The psychology of belief is something that is often neglected in traditional epistemology. This is unfortunate since just as there are relationships that must hold between our notions of truth, belief, and rational justification, these notions are in turn related to our understanding of rational inquiry. Insofar as scientific inquiry is supposed to provide knowledge, and knowledge is taken to involve rational justification, belief, and truth; scientific inquiry must provide inquirers with rational justification, and this rational justification must be such that in fact inspires belief of the appropriate scientific sort. Rational justification is useless unless it is in fact efficacious in establishing belief.

Belief is a psychological state over which we apparently have little control. While scientific inquiry should provide rational justification, and rational justification should track truth, and one's belief that it does track truth should lead to appropriate belief, whether a given set of methodological rules could in fact guide scientific inquiry depends on the actual psychology of the inquirers. If one believes that the truth link is satisfied by one's standard of scientific justification, then inquiry that provides scientific justification a belief would presumably lead to establishing the belief, but even this is not psychologically guaranteed. Sometimes one is led to doubt one's imagined methodological commitments because of the results they generate. This might be thought of as revealing other (meta-) methodological commitments. On this understanding of methodological commitments, inquirers will, of course, ultimately use precisely those commitments that they in fact have, whether or not they have rational justification for using them.

The point here is just that our notions of scientific belief, truth, rational justification, inquiry, and action are interrelated and mutually constraining. The suggestion is that one might use this fact to construct a theory of scientific knowledge by considering how each basic epistemic notion constrains the others. On this view no basic epistemic notion is primitive; rather, each is determined by the relationships it must bear to each other basic epistemic notion.

There is a rough analogy between this approach to scientific knowledge and how dynamical laws are investigated in modern physics. Experimental physicists try to get as much empirical evidence as possible concerning what physical quantities are conserved in fundamental particle interactions: energy, momentum, charge, lepton number, strangeness, etc. Each conserved quantity then corresponds to a symmetry that must be satisfied by the dynamics governing fundamental particle interactions. Each new symmetry that must be satisfied further constrains the dynamics. The hope is that a complete list of symmetries will uniquely determine the structure of our most basic physical laws. Just as one might take dynamical laws to be determined by the constraints that they must satisfy, one might take scientific knowledge to be determined by the constraints that must be satisfied between one's basic epistemic notions.

On a standard coherence account of knowledge, the rational justification or truth of a belief might be determined by the relationship the belief bears to empirical evidence or to other rationally justified or true beliefs; but here it is the notions of rational justification truth, belief, and inquiry that are themselves determined by the relationships they bear to each other. This is, of course, related to traditional epistemic coherence. Indeed, in a way, this is just a more thorough-going coherentism: one where one's account

of knowledge is itself a piece with one's other knowledge. I take this thorough-going coherentism to be one of the fundamental ideas in a particular brand of pragmatism: a brand represented perhaps most notably by C. S. Peirce and W. V. Quine. I also take this brand of pragmatism to embrace a thorough-going naturalism. I will focus here on Peirce's account of scientific knowledge.

II. Peirce's Account of Scientific Knowledge

It is often claimed that pragmatists define truth as usefulness: the formula is supposed to be that a proposition is true if and only if it is useful. But if this is the complete pragmatic story about truth, then it is a silly one. The most striking problem is that useful beliefs can nonetheless be false. This can be seen in the fact that showing that it is useful to believe something typically does nothing to provide rational justification for believing it. Another problem is that our pretheoretic notion of usefulness is itself too vague to be useful. A belief is useful if it satisfies some desired end, but the possible spectrum of such satisfactions is at least in principle mind-numbingly broad. Do we want to say that Plato's noble lie is in fact a noble truth because it plays an essential role in maintaining the political structure of the Republic? Do we want to insist that Zeus does in fact rule over Olympus in order to satisfy our religious yearnings? Do we want to aver that Murray Gell-Mann's eight-fold classification of fundamental particles reveals the true structure of the physical world because it reflects patterns from the *I Ching*? These are presumably not in the spectrum of *scientific* satisfactions since such satisfactions presumably have little to do with truth.[1]

Fortunately for pragmatism, usefulness need not be the whole truth about truth. More specifically, one need not take pragmatic truth to be *constituted* by usefulness. And C.S. Peirce did not. Rather, Peirce noted that a true proposition is useful because it is true, and it is learned to be true through its usefulness: a proposition's truth explains why it is useful and its usefulness in making accurate predictions, providing satisfactory explanations, and guiding action explains how it was found to be true.

Peirce's account of scientific knowledge is presented in the context of his analysis of the psychology of belief, doubt, inquiry, and success and failure action.

For Peirce, a state of belief is a comfortable state where one feels prepared to make judgments and to act. A state of doubt is an uncomfortable state where one feels unprepared for action. Indeed, doubt can render deliberate action impossible (Peirce 1877, 114).[2]

The irritation of doubt causes a struggle to attain a state of belief. This struggle to fix a belief is what Peirce calls inquiry. Inquiry is both initiated and terminated by the psychological state of the inquirer. Whether one is led to genuine inquiry is a function of how anxious one is concerning one's state of epistemic preparedness. If one finds oneself in a state of doubt concerning something that matters, then one enters into inquiry in order to fix a belief in the place of the doubt (Peirce 1877, 114).

The particular method of inquiry one employs is determined by one's current beliefs concerning how to best decide the truth concerning the matter at hand. Inquiry ends when one is comfortable, when one finds that a new belief has taken the place of doubt. What one ends up believing is what one takes as the truth for the purposes at hand. Inquiry does not continue until one has a true belief; rather, it continues until one

has a belief that one takes to be true. Since the fixation of belief ends the struggle that characterizes genuine inquiry, Peirce holds that the settlement of opinion is the sole aim of inquiry (Peirce 1877, 114-5).

In this sense, the process of inquiry is ubiquitous (Peirce 1877, 127-8). It is manifest when I tap my pocket to make sure I have the house keys before I lock the front door, and it is manifest when the apparent positions of stars are checked against the predictions of Einstein's general theory of relativity. And once those tests that are in fact taken to be required by inquirers are performed and passed, opinion is settled, and the settled opinion is what inquirers take as knowledge until doubt initiates further inquiry.

An inquirer would not require a belief that he did not doubt to pass any tests; but further, he would not know which test would might be relevant to inquiry without specific doubt. So without specific doubt, there would in fact be no inquiry because of the psychology of belief and there could be no inquiry since one would not know where to start. But, of course, we very often find ourselves well-stocked with specific doubts. These typically arise in light of unexpected failure in action.

A radical skeptic, or a philosopher, might charge that this rough psychological account of belief, doubt, inquiry, and action can have nothing to do with *knowledge* since, while it provides a rough description of how beliefs are formed and revised, it fails completely in explaining how they *ought* to be formed or revised. If an inquirer fixes a new belief, is comfortable, and inquiry stops, this can presumably mean nothing more than that the inquirer has fixed a new belief, is comfortable, and inquiry has stopped. Being comfortable enough to stop inquiry presumably does not, in itself, justify anything.

This immediately raises the traditional regress problem for justification: How does one ever fully justify one's actual standards of fixing belief?

If being comfortable enough to stop inquiry justified one's beliefs, then an inquirer who was so constituted that he never in fact doubted any of his beliefs would always have fully justified beliefs. While an inquirer will presumably believe that his beliefs are true (or probably true or probably approximately true), he presumably does not know that they are simply by dint of his believing them.[3] After all, if the inquirer had been born at a different time and place, if he had had different experiences or a different psychological constitution, he might well have believed something very different. One presumably wants an account of scientific knowledge that distinguishes between genuine scientific knowledge and mere opinion. More specifically, one might demand that an inquirer have some special sort of justification for a belief in order for it to count as scientific knowledge.

But perhaps this is where Peirce's brief analysis of the psychology of belief, doubt, inquiry, and action pays off. It is here, if Peirce has it right, where we understand how an inquirer might acknowledge that his beliefs are contingent on factors irrelevant to the truth of his beliefs without this having any effect whatsoever on his beliefs, commitments, or actions: the inquirer acknowledges the contingency, but the acknowledgment simply fails to have the psychological effect of inspiring any genuine doubt. The radical skeptic seeks to undermine belief by pointing out that inquirers fail to have rational justification for their beliefs. But the skeptic's charge that an inquirer only believes and does not have any ultimate, sure-fire justification for his beliefs is pragmatically (and truthfully) answered with "Yes, but I do *believe* it, and one of the

features of belief is that I take it to be true for the purposes at hand, and I consequently do not take further justification to be required. And insofar as your actions reveal your beliefs, you too are guilty of believing without the sort of justification that you claim to be necessary for rational belief." [4]

Peirce, like Hume before him, held that natural belief, regardless of whether or not it is somehow rationally justified, is enough to account for the deliberate actions of inquirers.[5] But, also like Hume, Peirce consequently risks being charged with having committed the naturalist's fallacy and mistaking explanation of one's beliefs for justification of one's beliefs. While, at some level of description, scientific knowledge can be nothing but opinion formed through natural psychological processes, this presumably cannot be the whole story. Bare, naturally fixed opinion (even if true) is presumably too cheap to count as scientific knowledge.

The skeptic then might be tempted to celebrate a victory. He has forced the (naturalistic, pragmatic) inquirer to admit that his current beliefs are formed from his past beliefs given his psychological constitution and experience. The inquirer recognizes that his beliefs are thus contingent on factors that may very well prove irrelevant to their truth. And the pragmatist further concedes that his standards of justification and method of fixing new beliefs are likewise contingent. But perhaps skeptical celebration is premature.

The inquirer concedes that he does not have any further justification for his current beliefs beyond their being the result of starting inquiry with the beliefs that he in fact started with and revising these on the basis of the best arguments and evidence he could find according to the best methodological principles he had, but what more could

he possibly do? Start from something that he did not actually believe? Use methodological principles that he doubts will lead to the truth? Not only would such moves be foolish, Peirce argues that they would be psychologically impossible: we have nothing but our current beliefs and commitments to guide inquiry, and we will apply them whenever we can.

[I]n point of fact, an inquiry to have that completely satisfactory result called demonstration, has only to start with propositions perfectly free from all actual doubt. If the premises are not in fact doubted at all, they cannot be more satisfactory than they are. (Peirce 1877, 115)

An inquirer can only be satisfied by beliefs formed in a way that he in fact takes to be appropriate. Once everyone believes something no further justification is needed for the belief to in fact be taken as knowledge: “When doubt ceases, mental action on the subject comes to an end; and, if it did go on, it would be without a purpose” (Peirce 1877, 115). Of course, it could happen that inquirers are brought to doubt something in the future that they do not doubt now; but if it is not doubted now, it is currently taken to be true for purposes at hand, without any reasons being given for believing it beyond those that in fact led to the belief being adopted and currently prevent doubt.

But then what is supposed to distinguish mere opinion from scientific knowledge for Peirce? Ultimately, I believe that it is the relationships it bears to past, current, and future inquiry that properly distinguishes scientific knowledge from opinion. While Peirce sometimes appears to dismiss the notion of rational justification altogether, there are at least two notions of justification that one might identify in his picture of scientific inquiry.

Diachronic justification for a belief involves providing the evidence and arguments that in fact led to the adoption of the belief. To give a diachronic justification for a belief is to give its genealogy. A diachronic justification tells the story of how a belief was fixed: it describes what the inquirer once believed and explains how the inquirer was led by his beliefs, empirical evidence, methodological commitments, and arguments considered to fix the belief.

Synchronic justification for a belief involves explaining how a belief is supported by other beliefs and commitments one currently holds. A synchronic justification might, for example, explain why one takes the belief to be reliable given what one believes about how the belief was formed. Such an explanation would be based on one's best understanding of one's reasoning and perceptual faculties and how these hook up with the world that one is investigating. Availing oneself of such explanations and believing what they say about the reliability of one's own beliefs is one of the ties that binds pragmatism and naturalism. Of course, the naturalist's explanation of the practical reliability of his scientific beliefs and commitments here depends on his science itself. But if one really has the relevant scientific beliefs and commitments, then why not use one's best account of the relationship between our beliefs and experience and the world to explain the believed reliability of one's scientific methodology? One would further expect that such an account might prove useful in suggesting ways to improve the reliability of the belief-fixing enterprise, and a good inquirer would be expected to exploit this knowledge.

Neither diachronic nor synchronic justification will satisfy the persistent skeptic. Just as the skeptic charges, all the scientific inquirer ever has is something he *takes* to be

true for reasons he *takes* to be good. The proffered justifications will only be recognized as justifications by inquirers who share the relevant beliefs and commitments of the inquirer providing them. An inquirer can explain how he came to hold the belief given what he once believed given, his method for fixing new beliefs, and the evidence he encountered. Or he can explain how his new belief meshes with his other beliefs to form a coherent account of the world and the success and failure of his actions. But, as the persistent skeptic would remind him, neither sort of justification provides the inquirer with unmediated access to the truth and both are contingent on the inquirers other beliefs and commitments, which themselves might be false.

On the other hand, the pragmatist finds it puzzling that anyone ever thought that one could do more than to report one's actual reasons for believing something when one has such reasons. How could one ever make any judgment that is not itself contingent on one's own nature, beliefs, and commitments? The pragmatist also finds it puzzling that anyone ever thought that more was somehow necessary. Why exactly would one want to satisfy the persistent skeptic's demands for ultimate justification when we usually find ourselves with sufficient knowledge to deliberate concerning even the most important of actions and to press forward with inquiry when we find our current knowledge lacking?

Peirce would agree with the persistent skeptic that neither sort of justification guarantees the truth of one's beliefs; but rather than undermining belief, which this observation itself clearly fails to do, Peirce takes it to demand a fallibilist attitude from rational inquirers. Since one can never be certain that one's beliefs are true, one must always be prepared to revise any or all of one's beliefs. And Peirce took this fallibilist attitude to be the fundamental unifying feature of the scientific method.

The fallibilist attitude involves acknowledging that it is always possible to be wrong about the nature of the world; and this, for Peirce, involves the commitment to there being real matters of fact concerning the real world that are independent of opinion. Insofar as the aim of science is to avoid error, science is constrained by matters of fact and is objective inquiry concerning the real. Indeed, Peirce believes that since science aims at the avoidance of error and since, as a scientist, he is committed to it being able to accomplish this aim, and since he takes the avoidance of error to involve knowing the truth, he believes that science will, in the limit of inquiry, provide the truth about the world. It is in this sense that a scientist might appear to be a sort of convergent realist. But for the fallibilist, we never in fact reach the limit of inquiry. (Peirce 1878, 138-9)

If providing a sure-fire guarantee of the truth of one's current beliefs is what it would take to answer the persistent skeptic, then a scientific inquirer will never do it. And as a fallibilist he will never try. But, if Peirce has the right analysis of the psychology of belief, the antics of the persistent skeptic will nonetheless eventually become irrelevant to inquiry as inquirers find themselves with settled beliefs. The skeptic asks inquirers to doubt their own settled beliefs. While the fallibilist will concede that his settled beliefs are not certain truths, to have a belief is not to doubt it. The persistent skeptic then is eventually ignored rather than answered. And one cannot help but wonder at the sincerity of his own claims to doubt.

III. Pragmatic epistemology and metaphysics

The question of how one ought to fix new beliefs when one is in a state of doubt can only be answered by careful reflection on one's own beliefs and commitments, and this is essentially what Peirce does as he takes his own current beliefs and commitments as evidence for what methodological beliefs and commitments that properly characterize natural science. He believes that scientific practice is grounded in the fallibilist attitude and that fallibilism only makes sense in the context of a realism that admits of objective matters of fact. Scientific methodology is thus characterized by its associated metaphysical commitments. It is the relationships that hold between belief, diachronic justification, and the real world that allow for the possibility of objective empirical knowledge. While this part of the argument has a transcendental flavor, it is, I believe for Peirce, thoroughly naturalistic.

Peirce's fallibilism requires that every feature of the scientific method is ultimately self-revisable in light of the empirical and explanatory successes and failures of the method in discovering the truth about the world. Since it is always subject to revision, all he can do is to describe the scientific method as he currently understands it and is committed to it. Given his commitment to fallibilism, there are no a priori, self-evident, certain first principles that could provide an absolute characterization of proper scientific inquiry. In this sense scientific knowledge can have no ultimate foundations.[6] One must start empirical inquiry with one's current methodological commitments. But again why did anyone ever think that one could do more than report what one in fact believes and describe the experience and arguments that one currently takes to justify the belief? And what more could one want than principles to which one is in fact fully committed?

It is possible that one begins inquiry with false beliefs or bad methodological commitments and is then led to mistaken conclusions about the nature of the world. But it is just this recognition that is reflected in the thorough-going fallibilism that characterizes natural science (Peirce 1877, 121).

While scientific inquiry is infinitely revisable, it is not arbitrary. Indeed, it always has precisely those methodological standards that the scientific community believes will lead to true belief. But since there are no certain first principles, scientific methodology must be self-correcting to reach the truth. And, for scientific inquiry to be objective, the beliefs of inquirers must be ultimately determined by some publicly accessible external permanency. To say that it is an external permanency is to say that it not constituted by the beliefs of inquirers but rather that it is what accounts for the success and failure in action that led to those beliefs. The external permanency is the world. It affects our senses according to regular laws. This, for Peirce, is the naturalistic content of the claim that we have epistemic access to the nature of the world by the empirical methods of natural science. Scientific aims at an objective description of the world as it is, as it determines our success and failure in action. That there is such an objective, empirically accessible world is, for Peirce, the “fundamental hypothesis” of the methodology of science (Peirce 1877, 120).

The persistent skeptic is unimpressed by such metaphysical musings. How does a scientific inquirer *know* that there are empirically accessible real things? For his part, Peirce does not even attempt to provide ultimate rational justification for his methodological and associated metaphysical beliefs. Rather he provides four reasons for why he holds them (Peirce 1877, 120-1). These four reasons provide naturalistic

explanations for his beliefs and commitments. His four explanations, I believe, naturally reduce to two. Both are psychological explanations.

When involved in scientific inquiry, Peirce notes that he has never in fact been led to doubt that he is discovering the nature of real things. The metaphysical commitments associated with the scientific method mesh with the actual practice and results of science. Using the method of natural science has led to success and the avoidance of error. The success of the scientific method and its ability to guide one away from error would be explained by the truth of the associated metaphysical commitments. But rather than claiming that this provides rational justification for realism, Peirce takes this to provide a naturalistic explanation for why he is never led to doubt the scientific method or the metaphysical foundations that he takes it to suppose. The explanatory coherence of the practice and the success of the scientific method and the associated metaphysics provides a psychological explanation for never in fact being led to doubt the relation between the method and the metaphysics.[7]

In addition to the explanation for why he has never been led to doubt his scientific commitments, Peirce provides a positive psychological explanation for why he finds himself with such commitments in the first place. He notes that everyone uses the scientific method when they know how to use it. For Peirce, this general principle of human nature is explained by the fact that the natural psychology of belief and doubt predisposes us to act as if there are objective matters of fact that can be determined through scientific inquiry. That doubt is an uncomfortable state and that it leads to empirical inquiry and that empirical inquiry is in fact efficacious in dispelling doubt provides a sense in which a positive commitment to objective truth is hardwired into our

psychology. If there were not objective matters of fact that could be determined through inquiry, then why would doubt be a matter of dissatisfaction that can be dispelled through inquiry? But, again, this is not meant to justify a commitment to there being objective matters of fact that can be determined through scientific inquiry; rather, this is a psychological explanation for why people tend to believe (insofar as their beliefs mesh with their behavioral dispositions) that there are objective matters of fact that can be determined through scientific inquiry. So if Peirce has the right psychology of belief, doubt, and inquiry, then our actions suggest an implicit psychological commitment to there being objective matters of fact that can be determined through empirical science.

If my reading here is right, nowhere does Peirce take himself to be providing an ultimate justification for his methodological and metaphysical commitments; rather he takes himself to be reporting what they are and explaining why they are compatible with his other beliefs and commitments. He reports his methodological commitment to fallibilism, explains the relationship between this commitment and his commitment to scientific realism, provides a psychological explanation for why he has never in fact been led to doubt these commitments in the practice of science, and explains why the psychology of belief and doubt predispose him, and anyone else who shares the psychology and the relationship between his actions and his beliefs and commitments, to believe that there are objective truths that can be discovered through scientific inquiry. He concludes:

These afford the explanation of my not doubting the method or the hypothesis which it supposes; and not having any doubt, nor believing that anyone else whom I could influence has, it would be the merest babble for me to say more about it.

If there be anybody with a living doubt upon the subject, let him consider it.

(Peirce 1877, 120-1)

Given his analysis of belief, doubt and inquiry, Peirce believes that everyone will find in his own actions and deliberations an implicit commitment to objective, accessible truth about the real world if one cares to look. But his own commitment to scientific inquiry is not contingent on other inquirers agreeing to there being objective truth that can be discovered by empirical science. Rather, his commitment to scientific inquiry and its objectivity is ultimately just something that he finds he has. He can explain why he has it, but he has the commitment before he provides the explanation.

While this chain of naturalistic explanations might look like a transcendental argument or some sort of coherence justification, Peirce clearly recognized these explanations do nothing to justify his beliefs to someone who fails to share his general naturalistic commitments. And all that Peirce can rely on in making his argument to the uncommitted is his belief that everyone's psychological make up is similar to his own and that others may thus be led to similar conclusions if they try to reconcile their actions with their beliefs and commitments.

While Peirce is ultimately uninterested in the arguments and opinions of people who fail to share his most general methodological commitments, he is supremely interested in the arguments and opinions of like-minded inquirers (Peirce 1877, 29). Along with basic empirical evidence one needs careful reasoning to arrive at the truth. A good scientist will consequently update his or her beliefs on the basis of good empirical evidence and good arguments. Taking such empirical evidence and compelling arguments seriously is part of what it means to be committed to scientific inquiry since this is the

best material at hand for progress toward the truth.[8] Peirce believes that scientists share a vision of scientific knowledge that ties together the notions of truth, inquiry, and the real world that science investigates:

[T]he followers of science are fully persuaded that the process of investigation, if only pushed far enough, will give one certain solution to every question to which they can be applied. ... [Various scientists] may at first obtain different results, but, as each perfects his method and his processes, the results will move steadily together toward a destined center. So with all scientific research. Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside themselves to one and the same conclusion. ... This great law is embodied in the conception of truth and reality. The opinion which is fated to be ultimately agreed to by all who investigate, is that which we mean by the truth, and the object represented in this opinion is the real. That is the way I would explain reality. (Peirce 1878, 138-9)

Such are the relationships that must hold between belief, truth, justification, success and failure in action, and the nature of the world on Peirce's pragmatic account of scientific inquiry.

IV. The diachronic structure of scientific knowledge

The notions of justification, belief, truth, success and failure action, and the nature of the world are clearly interdependent and mutually constraining on this model of scientific knowledge. No epistemic (nor metaphysical!) notion here is more fundamental

than another. And they can only be understood together by the role each plays in inquiry. On this account scientific knowledge has an essential diachronic structure.

Beliefs are reflected in action and updated on their success and failure in empirical prediction, explanation, and argument. When an inquirer encounters a piece of recalcitrant experience, his beliefs must be updated in order to account for the experience. He might change his background assumptions about the nature of the experiment that generated the experience. He might change her theoretical commitments. Or as Quine suggested, he might even plead hallucination or amend logical laws in order to account for his recalcitrant experience (Quine 1953).

There are rough general rules for how inquirers will revise their beliefs. Inquirers often try to balance maintaining systematicity and avoiding ad hoc revisions against producing the least possible disturbance to their core beliefs and commitments (Quine 1953, Quine and Ullian 1978, etc.). But precisely which beliefs are discarded, retained, and added will depend on a particular inquirer's methodological beliefs, commitments, and experience.[9] methodological commitments might in turn be revised in the context of their success and failure in inquiry. Just as one's other beliefs and commitments must answer to success and failure in action, one's methodological beliefs and commitments can be right or wrong insofar as they lead to right or wrong belief.

If an inquirer is to have any reason for revising his methodological commitments, it must presumably be by appeal to other methodological commitments that are somehow more firmly held. But there is nothing puzzling in this. Scientific inquiry has always sought for an ever more precise and thorough specification of those methodological commitments that reliably lead to truth. This aspect of scientific inquiry has led to

experimental controls, double-blind experiments, strict rules for treating empirical data in journal article and grant proposals, increasingly precise measuring devices, and increasingly sophisticated statistical tools. But, as good as they are, it is important not to mistake the status of our current methodological commitments: rather than being definitive of scientific inquiry once and for all, we fully expect that we will improve on our methodological commitments in currently unexpected ways as we better understand the world and our own nature as practitioners of science.

On Peirce's view scientific truth is what the scientific community would agree to when all of the empirical evidence is in and all relevant arguments have been given and evaluated, but it is important to note that the truth is not *constituted* by the consensus of the community. If it were, the scientific community could not be wrong when they all agree. But it is clearly possible for everyone to be wrong. And this fact is reflected in Peirce's uniform commitment to the fallibilist attitude. Armed with this commitment one will never be in a position to declare that one has reached the limit of scientific inquiry. Determining the truth then is a regulative goal for scientific inquiry, not the immediate end.

A truth relativist might want to say that that by determining truth, the consensus of the community likewise determines the nature of the real world and that this is how the truth manages to be truth about the world. But this completely misses one side of the relationship that must hold between the basic epistemological and metaphysical notions here. Just as the truth is not constituted by the limiting consensus of the scientific community, neither is the real world. The consensus that inquirers would reach in the limit of inquiry would be the result of their successes and failures in inquiry, and these

successes and failures are explained by the objective nature of the world. Again, if community consensus constituted truth, then it would be impossible for everyone to be wrong: that consensus does not constitute truth then can be seen in the fact that everyone can be wrong in the perfectly straightforward sense that their shared beliefs lead to such failure that they recognize that their beliefs are false. If there is anything noteworthy about this possibility, it is that it is so often realized. This possibility of failure is explained by the fact that even the scientific community's most firmly held and universally shared beliefs need not reflect matters of fact. It is the commitment to there being real matters of fact that leads to the scientific commitment to fallibilism, and it is the commitment to fallibilism that makes it impossible to mistake consensus for truth; but it is also the commitment to fallibilism that requires that one recognize the existence of objective matters of fact. On this view, *scientific* inquiry continues precisely as long as one holds that it is possible that everyone is mistaken. And it is in this sense that the fallibilistic attitude and the associated commitment to objective matters of fact lie at the heart of the scientific method.

The notions of truth and rational justification are also closely related on this account. Since diachronic justification involves explaining how one ended up with one's current beliefs and since the truth is reflected in the limiting beliefs of the inquirers, diachronic justification by its very nature tracks truth. But since we do not know what beliefs and commitments the future holds, we have no forward-looking sense of how well diachronic justification tracks truth. On the other hand, insofar as we have a backward-looking sense of scientific progress toward a better understanding of the world by comparing the various virtues and vices of our actual past theories against each other, we

have a backward-looking sense in which diachronic justification tracks scientific progress. And insofar as one can understand proximity to truth in terms of the avoidance of error and diachronic inquiry as the systematic elimination of error, one can see in a general way how diachronic justification tracks truth.

Rather than tracking some limiting notion of truth, synchronic justification tracks what one takes to be true. It appeals to one's current beliefs and commitments to explain why one takes his beliefs to be true. Synchronic justifications can be more or less interesting and persuasive depending on the specific evidence, beliefs, and systematic relationships to which one appeals and the degree to which these are shared by other inquirers. We can expect the standards we accept for rational justification to evolve just as our beliefs concerning the nature of the world to evolve since what we take to be true of the world and what we take as good standards of diachronic justification are interdependent. Indeed, one can expect that as one's beliefs and commitments are revised to avoid error, one will be able to provide better synchronic justifications.

While neither diachronic nor synchronic justification will answer to the persistent skeptic's demands, insofar as scientists agree in their beliefs and commitments, both diachronic and synchronic justifications are in fact taken to provide genuine normative justification in the practice of science.

Scientific inquiry aims at a reflective equilibrium between one's beliefs and methodological commitments, the beliefs and methodological commitments of the scientific community, and what is in fact true as reflected in success and failure in action. But insofar as one is committed to methodological fallibilism, one never expects this equilibrium to be satisfied.

V. Two problems dissolved

The infinite-regress problem for rational justification is in a way dissolved by a naturalistic account of the psychology of belief. Justification stops when there is in fact no real doubt. The persistent skeptic is not satisfied on his own terms; but insofar as he can always feign dissatisfaction and demand a justification for one's justification, there is good reason to suppose that the persistent skeptic can never be satisfied on his own terms.

[10] One might complain that the dissolution of this problem is not philosophically satisfactory, but I suspect that this just amounts to holding out for some special sort of epistemic certainty that would answer the persistent skeptic on his own terms. And, again, I do not believe that this is possible. All justification presupposes that something is granted. But, more importantly, there is a sense in which both the persistent skeptic and the philosopher are in fact already satisfied. Stopping when there is no real doubt for the purposes at hand is the standard everyone in fact uses in all deliberations concerning action. Even skeptics and philosophers apparently find this standard completely satisfactory for even their most important life-and-death decisions. It is this that makes any claimed dissatisfaction look disingenuous and appropriately ignored.

The reflectivity problem for a theory of knowledge concerns the epistemic status the theory grants itself. If each of a fallibilist's beliefs and commitments are revisable, then what is the status of his account of knowledge? What is the status of his commitment to fallibilism? Could his entire account be fundamentally mistaken? And if so, how ought he to revise his the very rules that explain how beliefs are to be revised?

Peirce presents his analysis of the relationship between belief, doubt, inquiry, and action as a piece of natural psychology.[11] In the construction of his account of knowledge he reveals his own empirical methodology for the fixation of belief and his naturalistic bias. But the resultant theory is fully believed and for what Peirce takes to be good empirical reasons. And for a fallibilist no theory can have a higher cognitive status than this.

The reflexivity problem is dissolved insofar as one applies precisely the same standards to the theory of knowledge that one applies to other objects of knowledge. On the current account knowledge is just current belief that is subject to revision in the light of empirical evidence and arguments and that answers to success and failure in action. A pragmatist would claim no other status for his theory of knowledge itself.

But there is still something puzzling here. Peirce's account of scientific knowledge rests on the possibility of genuine belief and commitment in the context of a thorough-going fallibilism. While there is certainly something curious about fully believing something and at the same time recognizing that one might be wrong, it is, it seems, nonetheless psychologically possible to be in this state. The recognition that one might be wrong comes from seeing how even our most cherished beliefs have been mistaken and from understanding how we form our beliefs and why we might thus be mistaken in our current beliefs. This recognition leads to a willingness to revise any belief in the light of relevant evidence and argument. But it typically in fact does little to undermine our current beliefs or prevent deliberate action. And this is apparently just a natural psychological fact.

A fallibilistic attitude can apparently even be maintained when one believes something to be necessarily true. One might not be able to imagine how the belief could

possibly be false, it might be literally inconceivable, but once one recognizes that this is all that characterizes judgments of necessity, it is easy to allow for the possibility of being wrong. What is conceivable is itself clearly contingent. Here possibility is psychological possibility. Of course, one does not believe that one's beliefs are wrong, especially not the necessary ones, but to be a fallibilist is to recognize that inquiry is open-ended, that we can even make mistakes in our most cherished logical, mathematical, and philosophical judgments. It is to settle for contingent judgments of necessity.

The psychology of belief then allows for a naturalistic explanation for how one can be simultaneously sectarian and ecumenical concerning one's methodological beliefs and practices. One really does believe what one believes, but one can at the same time recognize that believing something is not constitutive of its truth and that our beliefs are contingent on factors that are obviously irrelevant to their truth. There is nothing contradictory in holding to both sides here, especially since they are both clearly right. And it is the psychological possibility of both believing what one believes and believing that these beliefs could be false that allows for the possibility of scientific inquiry insofar as it requires belief and is grounded in a thorough-going fallibilism.

The naturalistic pragmatist's account of knowledge is itself part of his beliefs and commitments, and as such, it could clearly be wrong. But if it is, then the pragmatist who is committed to the account has nothing to say about how it might be wrong or how it should be revised. Indeed, if one is fully committed to the account, it may be inconceivable exactly how it could be wrong. This does not mean that one cannot be mistaken, it just means that one cannot see how one *could* be mistaken. In this sense, this account of scientific knowledge might be understood to be necessarily true if one is to

provide an account of scientific knowledge at all. The account of knowledge might still be fundamentally wrong, but one cannot imagine how.

VI. Progress and the transcendental character of scientific knowledge

In his attack on naturalized epistemology, and Quine's naturalism in particular, Putnam argues that an adequate account of knowledge must capture both its immanent and transcendent character (Putnam 1983). Our standards of rational belief are immanent in that they are *our* standards. They are the product of our experience, psychology, physiology, history, and culture. They are, consequently, fundamentally contingent. But we also take our standards of rational belief to be transcendent. They are supposed to be more than just what we happen to be committed to right now. They are supposed to provide genuinely good reasons for the beliefs that we have and be able to act as general rules to guide future inquiry. They are supposed to make the difference between contingent opinion, which is too cheap, and real knowledge, the stuff that one somehow has a right to believe. Putnam complains that while naturalized succeeds in capturing the immanent nature of knowledge, it fails to capture its transcendental nature.[12]

But perhaps the naturalistic epistemologist has a satisfactory response. For the naturalistic epistemologist it is certainly true that all of one's beliefs are both immanent and contingent. Indeed, insofar as one is interested in providing a naturalistic account of belief, one would presumably be interested in revealing the many ways in which even one's most cherished beliefs are contingent on one's immanent psychology and environment. But even so, this does not mean that an inquirer's beliefs and commitments are merely contingent. Perhaps most significantly, an inquirer's own beliefs and

commitments have a special status for the inquirer himself. They are, after all, his beliefs and methodological commitments. His beliefs are taken by him to be true, and his methodological commitments are taken by him to be those rules that will in fact lead to successful future inquiry. From the perspective of the inquirer then his beliefs and commitments are taken to have a transcendent character. Of course, his perspective is both immanent and contingent, so his judgments of transcendence are themselves immanent and contingent. But for the naturalist all of our judgments are both immanent in that they are ours and contingent in that they depend on the natural details of ourselves and our environment. On this view, it is only possible to make immanent judgments of transcendence and contingent judgments of necessity, and it is sophistry and illusion to suppose that one can do more.

The naturalist's recognition that all judgments are both immanent and contingent carries with it an essential risk and a corresponding responsibility. Since one always might be wrong, one can never be certain. And this means that if one wishes to avoid error, one must always be prepared to revise even one's most cherished beliefs and commitments. The naturalist's recognition then requires the methodological commitment to fallibilism.

As we have seen, methodological fallibilism is itself a transcendental notion. The possibility of being wrong involves an implicit commitment to a standard of truth that transcends one's own beliefs and commitments. And the possibility of everyone being wrong involves an implicit commitment to a standard of truth that transcends the consensus of the community of inquirers. The naturalist does not claim immediate epistemic access to objective truth. But insofar as he is committed to epistemic progress

by revising his beliefs in the light of their successes and failures, he is committed to a transcendent standard for inquiry. It is the nature of the world that explains success and failure. Revision of our beliefs and commitments in the light of our experience and best arguments is what guarantees that the world is, in just this sense, reflected in our beliefs. The truth about the nature of the world provides a vague, transcendental regulative ideal for inquiry: vague because we do not know how our beliefs and commitments will in fact change and transcendental because truth is not constituted by opinion. While it might be comforting to know that there is a transcendental regulative standard for inquiry, since we do not have epistemic access to the results of future inquiry, this vague forward-looking standard provides no guidance beyond requiring that we eliminate error in our beliefs and commitments whenever we find it. One way to see its transcendental nature is to note that methodological fallibilism involves an implicit commitment to scientific progress.

There are several concrete, backward-looking notions of scientific progress to which one might appeal. Insofar as we can specify the content of our scientific theories, we can say exactly how our current theories have succeeded where our past theories failed. There are different backward-looking notions of progress for each of the relative theoretical virtues one might consider. One might, for example, praise Copernican astronomy for giving up the Ptolemaic earth-centered system, yet recognize that its empirical predictions were not as accurate. Or one might concede some explanatory virtues to Cartesian over Newtonian mechanics yet recognize the ad hoc nature of the former theory. Just as there are many ways of characterizing scientific progress, there is no single, backward-looking notion of truth. There are many ways that one might compare theoretical options, and one might associate a sense of proximity to the truth

with any of these. If a particular virtue comes in measurable degrees, one might even impose a metric that would allow one to judge how close rival theories are to each other or to some other standard with respect to the virtue.

Scientific progress is often characterized in terms of the empirical accuracy and scope of our theories. One might compare the empirical predictions of our theories against our experience and judge one theory to be better than another if its predictions are more accurate. Or one might compare the relative domains in which rival theories make accurate predictions and favor the broader theory. Or one might simply compare our current best theories against our best empirical evidence in order to *judge in this limited sense* their absolute accuracy. Given our track-record, one might expect science to develop increasingly accurate empirical theories. This sense of absolute progress in empirical accuracy provides an historically important notion of approximate truth: a theory is approximately true insofar as it agrees with our current best empirical evidence. But like the corresponding sense of progress, this positivistic notion of truth is relative to what we take as our current best empirical evidence and how we understand it. This notion of approximate truth is also relatively weak in that it fails to take into account such things as the explanatory virtues of rival theories.

One would obviously like to be able to compare our current theories to the truth, to an account of the world that is completely satisfactory in every predictive and explanatory detail. But, not only is there no single, canonical way to carry out such a comparison if we had unmediated access to the truth, we don't. Since we cannot compare our current theories in any way against an account we do not have, we are left unable to make concrete forward-looking judgments concerning how our theories are

approximately true as descriptions of the world. Perhaps our vague forward-looking notion of scientific progress and approximate truth is fully captured in our commitment to revise our scientific beliefs in the light of future empirical evidence and relevant arguments. This conclusion might be disappointing, but it is not surprising. That we can make no concrete forward-looking judgments of approximate truth might be thought of as the price we pay here for wanting more from a true theory than empirical adequacy. While we lack a concrete notion of approximate truth, we do have a vague forward-looking notion of approximate truth. While we do not yet know exactly how, we nonetheless believe that our future theories will be better than our current theories in significant ways, that there are mistakes we are making now that we will not make in the future. This forward-looking notion of progress is an important component in our account of scientific knowledge. But this is only a vague notion of progress since we lack a concrete comparison of our current theories against our future theories. Since we do not know the ways in which our current theories will fail and how we will address these failure, one cannot say the sense in which our current theories are true or false. We often believe that we are on the verge of getting complete and accurate scientific theories. This appeal to our vague forward-looking notions of scientific progress and approximate truth is certainly one of the psychological forces that drive natural science. But if we ever did give up the possibility that we might be wrong in our best theoretical commitments, then scientific inquiry would end.

Not being certain that one's current theories are true or even approximately true in a concrete forward-looking sense need not, and typically does not, in itself generate doubt concerning the truth of one's current scientific beliefs and commitments. And while it is

the vague forward-looking notion of the truth about the real world that explains how one's judgments could be wrong, since we have no access to the truth except through scientific inquiry, our current scientific beliefs are our best rational beliefs concerning the nature of the world.

It is our commitment to the open-ended nature of inquiry, to taking into account all relevant future evidence and arguments that are relevant, that makes our scientific beliefs and commitments more than just contingent opinion. Our scientific beliefs and commitments have an essentially diachronic nature. They were forged in the context of our best past evidence and arguments where we took into account everything that we took to be relevant to their truth, and they both shape and answer to our best future evidence and arguments.

VII. Unsolved problems and conclusion

Perhaps the most serious problem with the present account of scientific knowledge concerns the proper cognitive status of our best scientific theories. So far, we have supposed with Peirce that scientific theories are believed and that belief is a property that either obtains or does not. But, as noted earlier, since scientific justification comes in degrees, one would expect that scientific beliefs should come in degrees as well. When our best physical theories predict that we will find a new sort of fundamental particle with such and such properties, we typically neither fully believe the prediction nor fully disbelieve it. One would presumably want an account of scientific knowledge that captures such degrees of belief and explains how rational degrees of belief should be

updated in the light of new evidence. But faithfully representing the cognitive status of our best scientific theories requires more than just capturing degrees of belief

We typically do not take even our best, most detailed, and precise scientific theories to be even probably true. Indeed, there are almost certainly mistakes in the details of our most ambitious theories. This is no doubt why it is suggested that we take our best theories to be neither true, nor probably true, but only probably approximately true. But, again, there is no single, canonical *backward*-looking notion of approximate truth, and there is no concrete forward-looking sense such that we can judge our best current theories to be approximately true as descriptions of the world.

Laudan (1981 and 1987) has argued that scientific progress be characterized by concrete progress in problem-solving power rather than by any notion of our increased proximity to truth. This provides an option for the pragmatist who is uncomfortable with talk of scientific truth. But I do not think that one needs to get rid of the notion of truth in order to account for scientific knowledge. After all, there are presumably parts of science that we believe to be true. Electrons are negatively charged. The earth is roughly spherical. Table salt is sodium chloride. DNA is involved in the propagation of an organism's physical traits to its progeny. Many scientific claims have the same right to be considered true as any other commonsense claims about the nature of the world. While one might wonder just what sort of truths these are, we have to puzzle over the sincerity of anyone who has serious doubts concerning such claims, and any notion of truth that fails to make such claims true is probably too fancy to be itself useful.

This said, some of the most significant scientific postulates and theories clearly fail to have the cognitive status of simply being believed to be true. The standard collapse formulation of quantum mechanics is a salient example. It is clearly an important part of our physical knowledge, yet it is also just as clearly false. And it is unclear how it might be fixed.[13]

Rather than being believed to be true, probably true, or even probably approximately true, it seems that the open-ended nature of scientific inquiry is often baked into the cognitive status of our best theories. Our best theories provide the best predictions, explanations, and understandings of the world that we have. They have been forged in the context of the best empirical evidence and arguments that we have taken to be relevant to understanding the world. And we use them to guide our most important actions and to suggest future inquiry. But rather than believe that they are true or probably true or approximately true, we often take them to be true for the purposes at hand or useful toward a specific end (which may include furthering scientific inquiry itself). Even when we take a theory to provide our best understanding of the world or to provide good explanations, this need not involve taking the theory to be true, probably true, or probably approximately true.

Scientific inquiry provides us with our best understanding of the world with our best empirical predictions. The understanding and predictions are arrived at through our best and most careful empirical investigations, and they take into account all of the relevant arguments and empirical evidence that we currently have. And this, it seems to me, is the general cognitive status of our best scientific theories.

On this account, scientific knowledge can be understood in terms of the relationships that hold between scientific belief, inquiry and justification, and the nature of the world as reflected in our success and failure in action. Scientific inquiry begins with those beliefs we in fact hold about how to best investigate the nature of the world, and it leads toward truth through the systematic elimination of error. Truth is truth about the nature of the real world. That there are matters of fact about the world explains why no consensus constitutes truth and it explains science's associated methodological commitment to fallibilism. The nature of the world explains our specific success and failure in action. And our success and failure in action provides the standard for revising our beliefs, the methodological commitments of science, and the standards of scientific justification. No basic epistemic notion is more fundamental than another; rather, each is constrained by the systematic relationships that must hold between them in a naturalistic account of scientific knowledge.

Endnotes

1. William James included such satisfactions among those that could bear on the avoidance of error and the truth of a belief. See, for example, James's 1896 (1956) essay "The Will to Believe" sections IX and X. But scientific knowledge is presumably characterized by a much narrower spectrum of possible satisfactions, though perhaps broader than many philosophers of science have suggested. Perice sought to distance his own views on knowledge from James's.

2. The uncomfortable feeling of doubt can be approximated by imagining being half-way across a rope bridge over an abyss and noticing that the ropes are deeply frayed so that one genuinely doubts the safety of the bridge and it is unclear whether one should move forward, back, or stand still. If the stakes are sufficiently high, such doubt might generate complete paralysis. That believe and doubt are a function of the stakes is something that Peirce himself did not pursue. If he had, he might have developed a notion of degree of belief corresponding to preparedness for action at various stakes.

3. That mere opinion, even if true, is too cheap to count as knowledge is arguably the insight on which epistemology was founded. In the *Theaetetus*, after pointing out that the members of a jury might judge correctly but without knowledge, Socrates concludes that true belief and knowledge must be different things. Theaetetus replies that he had already heard someone make this distinction then argue that knowledge requires true belief with the addition of *logos*. Socrates seems to approve of this suggestion. (Hamilton and Cairns, eds. 1987, 908). One would presumably want one's account of knowledge to distinguish between opinion and knowledge since without this, one would arguably not have an account of *knowledge* as the term is in fact used.

4. This is the heart of Peirce's famous attack on Cartesian doubt (Peirce 1868a, 11-27 and 1868b, 28-30).

5. For his part Hume held that reason was too weak to guide action. Fortunately for us, he explained, we find ourselves endowed with natural instincts and habits that lead us to

successful action (1777; 1994, 43-5). This move from the failure of rationalism to a pragmatic naturalism is mirrored in Peirce's early arguments against Cartesian doubt. Hume seems to have missed the point that our faculties of reason, even those involved in judgments concerning the relations of ideas, are not different in kind from our other natural cognitive faculties (like those involved in judgments concerning matters of fact). For his part, Peirce, seems to have embraced a thorough-going naturalism with respect to all our faculties of inquiry, a move that Quine later follows (Quine 1969).

6. In this sense even the commitment to fallibilism and the associated metaphysical commitments are negotiable. Consistency requires that one be a fallibilist with respect to one's fallibilism. If one were to convince oneself that one had found certain foundations for scientific knowledge and were to consequently relinquish one's commitment to fallibilism, this would forever change the fundamental nature of inquiry for anyone who shared the new methodological commitments. This is one way to explain what happened to Descartes in the context of Peirce's own account of scientific knowledge. But, for his part, Peirce cannot imagine himself with certain epistemic foundations.

7. While this bears a superficial similarity to more recent arguments for realism from the success of science, there is a critical difference (for a review of some such arguments for realism see Fine 1984). If I am right, Peirce is providing a psychological explanation for why he has not yet been led to doubt his scientific realism, not an argument or justification for the position. And this is empirical, not transcendental, psychology. There is, however, a problem. Given the skeptical induction on the failure of science to

get the whole truth and nothing but the truth, it is not as clear to us the sense in which the practice and results of science suggest the existence of an empirical accessible objective world. On the other hand, however our theories evolve, they are never incompatible with there being such a world. And, as explained in his positive account of our metaphysical commitments, requiring an explanation of our success and failure in action and that our beliefs are consequently not constitutive of truth, suggests that there is a world.

8. It is only because scientists take each other's positions seriously that quick, rational consensus is possible in the scientific community. This practice is justified if it has proven a reliable or fruitful way of fixing the sort of scientific belief for which it is used. See Franklin (1990) and Barrett (1996) for Bayesian flavored reflections on this practice.

9. This general phenomenon can be seen in the current debate over surreal trajectories in Bohm's theory. One camp of physicists takes the existence of surreal trajectories in the theory to entail that Bohm's theory is false while another camp takes this feature of the theory to entail that our measuring devices do not work the way that one might have initially thought. Since both camps agree that Bohm's theory makes the right empirical predictions, there is no canonical way to decide the issue. How it is ultimately decided will probably depend on how successful the Bohmians are at extending their theory to cover a broader domain of quantum phenomena than the basic theory can cover. Finding some way to reconcile the theory with special relativity, for example, would be a significant advance. If they are successful enough on other fronts, then the physics community might be willing to accept the counterintuitive surreal trajectories predicted

by the theory. See Barrett (2000) for a detailed description of this debate and the concrete sense in which it illustrates the theory-ladenness of observation.

10. Quine (1969) and others have argued that the doctrinal or normative project in epistemology be abandoned. If the aim is to find, once and for all, the proper, self-evident, and certain foundations of knowledge, then this project should certainly be abandoned. But if the project is to describe what we in fact take to be reliable evidence and arguments for our belief and explain why, then this is something that one might expect to be complicated but not impossible. For the fallibilist, the details of such stories are, of course, revisable.

11. Just as Quine (1969) would later have it.

12. Putnam's (1983) notions of immanence and transcendence roughly correspond to the traditional notions of explanation and justification. And his charge that naturalized epistemology fails to capture the transcendent character of knowledge roughly corresponds to the naturalistic fallacy.

13. In the case of quantum mechanics we have very good reason to suppose that the standard collapse theory is false, and we have nothing more probable to take its place.

See Barrett (1999) for details concerning this problem.

Bibliography

Barrett, J. A. (2000) "The Persistence of Memory: Surreal Trajectories in Bohm's Theory", *Philosophy of Science* V67(N4):680-703.

(1999) *The Quantum Mechanics of Minds and Worlds*, Oxford ; New York : Oxford University Press.

(1996) "Oracles, Aesthetics, and Bayesian Consensus", *Philosophy of Science* V63(N3):S273-S280.

Carnap, R. (1950) "Empiricism, Semantics, and Ontology", *Revue Internationale de Philosophie* 4:20-40.

Fine, A. (1984) "The Natural Ontological Attitude", in J. Leplin ed., *Scientific Realism*, Berkeley: University of California Press (1984), 83-107.

Forster, P. D. (1994) "The Limits of Pragmatic Realism (C. S. Peirce and the Limit Theory of Truth)", *Philosophy Today* V38(N3):243-258.

Franklin, Allan (1990) *Experiment, Right or Wrong*, Cambridge ; New York :
Cambridge University Press.

Hamilton, Edith and Huntington Cairns, eds. (1987) *Plato: The Collected Dialogues*,
Princeton, NJ: Princeton University Press.

Houser, Nathan and Christian Kloesel, eds. (1992) *The Essential Peirce*, Bloomington;
Indianapolis: Indiana University Press.

Hume, David (1994) *Enquiries Concerning Human Understanding and Concerning the
Principles of Morals*, Oxford: Clarendon Press, reprinted from the posthumous
1777 edition.

James, William (1896) "The Will to Believe", originally published in *New World* 5 (June
1896): 327-47, reprinted in James, W. (1956) *The Will to Believe and Other
Essays*, New York: Dover Publications Inc.

Laudan L. (1990a) "Aim-less Epistemology", *Studies in History and Philosophy of
Science* V21(N2):315-322.

(1990b) "Normative Naturalism", *Philosophy of Science* V57(N1):44-59.

Lloyd, E. A. (1997) "Feyerabend, Mill, and Pluralism", *Philosophy of Science*

V64(N4):S396-S407.

Peirce, Charles Sanders (1868a) “Questions Concerning Certain Faculties Claimed for Man”, originally published in *Journal of Speculative Philosophy* 2 (1868):103-14; quoted here from Houser and Kloesel, eds. 1992, pp. 11-27.

(1868b) “Some Consequences of Four Incapacities”, originally published in *Journal of Speculative Philosophy* 2 (1868):140-57; in Houser and Kloesel, eds. 1992, pp. 28-55.

(1877) “The Fixation of Belief”, originally published in *Popular Science Monthly* 12 (November 1877):1-15; in Houser and Kloesel, eds. 1992, pp. 109-23.

(1878) “How to Make Our Ideas Clear”, originally published in *Popular Science Monthly* 12 (January 1878):286-302; in Houser and Kloesel, eds. 1992, pp. 124-41.

(1892) “The Doctrine of Necessity Examined”, originally published in *The Monist* 2 (April 1892): 321-37; in Houser and Kloesel, eds. 1992, pp. 298-311.

Putnam, Hilary (1983) “Why Reason Cannot be Naturalized”, from *Realism and Reason*, volume 3, *Philosophical Papers*, Cambridge: Cambridge University Press (1983), 229-47.

Quine, W. V. (1969) "Epistemology Naturalized", from *Ontological Relativity and Other Essays*, New York: Columbia University.

(1953) "Two Dogmas of Empiricism", from *From a Logical Point of View*, Cambridge, MA: Harvard University Press.

Quine, W. V. and J. S. Ullian (1978) *The Web of Belief*, 2d ed. New York : Random House.