

January 2012

The View from the Armchair: Responding to Kornblith's Alternative to Armchair Philosophy

Anthony Bryson
South Central College

David Alexander
Huntington University

Follow this and additional works at: <https://commons.pacificu.edu/eip>

 Part of the [Philosophy Commons](#)

Recommended Citation

Bryson, Anthony and Alexander, David (2012) "The View from the Armchair: Responding to Kornblith's Alternative to Armchair Philosophy," *Essays in Philosophy*: Vol. 13: Iss. 1, Article 10.

The View from the Armchair: Responding to Kornblith's Alternative to Armchair Philosophy

Anthony Bryson, South Central College

David Alexander, Huntington University

Published online: 30 January 2012

© Anthony Bryson, David Alexander 2012

Abstract

In the last two decades, the greatest threat to armchair philosophy has been the natural kinds approach. On this view, philosophic theorizing should not be obsessed with the *ideas* of justice, goodness, and truth but should look outward to the world of *objects* to find these things. And if these things happen to be natural kinds, like kinds of rock or fish for instance, then clearly we should reject the armchair for the lab. The philosopher should leave the office and join the scientist out in the field. Philosophy should become a species of science. We attempt to defend traditional/armchair philosophy by examining Hilary Kornblith's naturalistic methodological approach to epistemology. Among other things, we argue that Kornblith's approach leads to some surprising, undesirable results (at least undesirable to the naturalist), one of which is that Kornblith cannot discount epistemic internalism as a viable contender in the search for the nature of knowledge. His methodology actually requires that we take epistemic internalism seriously.

In the last two decades, the greatest threat to armchair philosophy has been the natural kinds approach. On this view, philosophic theorizing should not be obsessed with the *ideas* of justice, goodness, and truth but should look outward to the world of physical *objects* to find these things. And if these things happen to be natural kinds, like kinds of rock or fish for instance, then clearly the philosopher should leave the office and join the scientist out in the field. Philosophy should become a species of science.

In what follows, we will address the methodological naturalism found in the work of Hilary Kornblith. In *Knowledge and Its Place in Nature* (2002) and in an earlier piece "The Role of Intuition in Philosophical Inquiry: An Account with No Unnatural

Ingredients” (1998) Kornblith challenges the more traditional conception of philosophy. His approach to philosophical issues is intriguing and has the salubrious effect of making traditional epistemologists more self-aware of their methods. Yet, if Kornblith is right and the search for knowledge ought to be modeled on the search for natural kinds, then the investigations of armchair philosophers like ourselves misses wide of the mark.¹

To Kornblith’s approach to philosophy, we apply the label “extra-mentalism,” for he proposes that we discover the nature of philosophical objects by looking outward and not at our representations of them:

On my view, the subject matter of ethics is the right and the good, not our concepts of them. The subject matter of philosophy of mind is the mind itself, not our concept of it. And the subject matter of epistemology is knowledge itself, not our concept of knowledge. (2002, 1)²

We must admit a feeling of kinship toward Kornblith at least at this stage. If math and logic can be done from the armchair and seem *prima facie* to be about something other than language or thought, why can’t philosophy be as well? Why can’t extra-mentalism be true for armchair philosophy? But this is as far as our sympathies coincide, for Kornblith’s brand of extra-mentalism would have us forsake the armchair for empirical investigation.

Armchair philosophy has several different variants. The most popular is probably mentalism. Mentalists believe that mental properties, like concepts, are the targets of philosophical analysis. We should not look to the mind independent world but to our representations to learn the conditions for such things as knowledge or free will. In conjunction with this view there is often the unstated assumption that our concepts determine the structure of philosophical phenomena. This has the welcome result that we can discover what free action is in a world where no one acts freely. Genuine philosophical analysis doesn’t require the existence of actual free acts or actually good things.³ Indeed, some would say that we need to get a good way along in the analysis before we can get strong justification for thinking the things we’re talking about exist.⁴ Those sympathetic to epistemic skepticism probably believe something like this.⁵ So no matter what world we occupy, we can still give a correct analysis of free action or the morally right.

This view of conceptual analysis according to which our representations determine the structure of philosophical objects is attractive to the armchair bound philosopher for another reason. This kind of conceptual analyst can explain why intuitions (usually the most important source of data for the armchair philosopher) have evidential weight. Concepts not only determine the nature of philosophical phenomena, they engender data carrying intuitions. If however philosophical objects have an ontologically prior existence to our concepts, then the conceptual analyst must show how our concepts mirror the structure of the world (assuming intuitions come from concept possession). Though this

view of conceptual analysis is the one we're partial to, many see any attempt by the armchair philosopher to connect our concepts with these allegedly prior existing objects as a forlorn undertaking bound to end in failure. This is primarily because they're suspicious of any kind of a priori knowledge that extends to things outside the mind.

For Kornblith, we don't need to consult our concepts to get at the nature of things like knowledge or free will, for if they exist, they exist out there in the physical world to be discovered. Just as scientists amend and emend their concepts in the face of scientific discovery, philosophers should do the same with theirs. In fact, if Kornblith is right, we have another reason to give up the introspective gaze: our concept of knowledge can misrepresent what knowledge is. Our epistemic concepts, and not just our analyses of them, can be more or less correct depending on how well they accurately characterize a natural kind.⁶ We need not worry then (as some have) whether our concept of "justice" is the same as Plato's, for this directs us away from the true object of study—justice as a natural kind.

But how could one possibly discover the nature of knowledge empirically? Can we point to knowledge? Does knowledge look a certain way, sound a certain way, or taste a certain way? Can we stumble over knowledge? We shall now see what Kornblith has to say on these matters.

Kornblith's Proposal

We title this section "Kornblith's Proposal" because, in *Knowledge and Its Place in Nature*, he makes several assumptions (which he may leave others to argue for) when sketching his new approach to philosophy. He assumes for the sake of argument a causal/historical account of reference.⁷ Moreover, he probably assumes a thoroughgoing physicalism, a view often hard to pin down, but that can presently be defined as the belief that all that exists is physical or is the sort of thing that enters into causal relations nomologically understood. Combining these with the view that knowledge is either a natural kind or nothing at all, we realize that epistemic internalists or traditional foundationalists can't hope to see Kornblith vindicated and then use his approach to prove their views in epistemology. His method will restrict the number of legitimate theoretical outcomes. Prima facie, this may look unfair as though Kornblith is, in a sense, rigging the debate between epistemic internalists and externalists. He takes what has often been considered a foundational branch of philosophy—epistemology—and uses other areas of philosophy to reach a certain epistemological outcome—in particular some major commitments in metaphysics and philosophy of language.

Alvin Goldman's earlier epistemology (and to a lesser degree his later epistemology) employs a less restrictive method.⁸ Though Goldman is a reliabilist like Kornblith and as a reliabilist would have us empirically determine which belief-forming processes we use and whether they're reliable, he still does meta-epistemology from the armchair. Only after he completes the analysis and learns that justification involves reliable causal mechanisms does he then recommend that we turn to cognitive science. Thus an internalist could agree with the way Goldman does philosophy while disagreeing with his analysis. This is not true of Kornblith's naturalistic methodology. We pretty much know using his approach that certain internalist theories won't be viable options (especially those that posit *sui generis* relations like acquaintance). This is a point we will press later.

For the armchair philosopher, the data of philosophy is mostly intuition. Examples of knowledge are brought before the mind via intuitions about specific cases and those examples are examined for crucial common properties. Some even argue that intuitions could be about principles and not just specific cases. Utilitarians tend to stress the intuitiveness of certain moral principles and those of a skeptical bent tend to do the same in epistemology.⁹ Kornblith chooses to keep many of the same terms, but gives them a different characterization. He is happy to say that we pick out samples of knowledge using intuitions. But for him, intuitions are plain old empirical judgments. They are not *a priori* in any traditional sense:

When we appeal to our intuitions about knowledge, we make salient certain instances of the phenomenon that need to be accounted for, and that these are genuine instances of knowledge is simply obvious, at least if our examples are well chosen. What we are doing, as I see it is much like the rock collector who gathers samples of some interesting kind of stone for the purpose of figuring out what it is that the samples have in common. We begin often enough, with obvious cases, even if we do not yet understand what it is that provides the theoretical unity to the kind we wish to examine. (Kornblith 2002, 11)

A page later he writes:

[O]n the account I favor, these resulting judgments are no more *a priori* than the rock collector's judgments that if he were to find a rock meeting certain conditions, it would (or would not) count as a sample of a given kind. All such judgments, however obvious, are *a posteriori*, and we may view the appeal to intuition in philosophical cases in a similar manner. (2002, 12)

Thus as we understand Kornblith, we find samples of knowledge by initially relying on intuitions about particular cases. These intuitions amount to ordinary common sense empirical judgments no different than the judgment that something is a rock or gold. Once we have enough cases collected, we can then empirically investigate them for a theoretical unity. The process will be a matter of trial and error just as a scientific inquiry into rocks would be. We may find that we have to discard some of our initial samples since they fail to belong to the kind in question, or that properties we initially took to be unifying are not.

Consider the way he describes investigating rocks, a process he later uses to describe epistemic investigation:

What should we say about the rock collector's judgments at early stages of investigation, i.e. prior to any deep theoretical understanding of the features that make his samples samples of a given kind? Such judgments are, of course, corrigible, and they will change with the progress of theory. What seemed to be a clear case of a given kind in the absence of theoretical understanding may come to be a paradigm case of some different kind once the phenomena are better understood. (2002, 13)

Initially, the rock collectors might think there are some characteristics, like smell, which distinguish kinds of rocks or rocks from non-rocks only later to find they are mistaken.

Once we get through the initial stages, our competency and understanding of what sorts of samples to look for increases. Our judgments about which cases constitute knowledge grow epistemically stronger as we move along in the process and acquire better samples to empirically investigate.

The judgments of rock collectors at early stages of investigation are substantially inferior, epistemically speaking to those at later stages, when theoretical understanding is further advanced. We should not say that initial judgments are of no evidential value, for were this the case progress in theory would be impossible. Our untutored judgments must have some purchase on the phenomenon under investigation; but, that said, it must also be acknowledged that judgment guided by accurate background theory is far superior to the intuitions of the naïve. (Kornblith 2002, 14)¹⁰

Our goal then, according to Kornblith, is to discover a natural theoretical unity among the various samples, to discover their residence in the category of a natural kind—one that can do explanatory work in a scientific view of the world. Not until later does Kornblith tell us how he understands a natural kind:

Following Richard Boyd, I take natural kinds to be homeostatically clustered properties, properties that are mutually supporting and reinforcing in the face of external change...the properties that are homeostatically clustered play a significant causal role in producing such a wide range of associated properties, and in thereby explaining the kind's characteristic interactions. It is for this reason too that natural kinds feature so prominently in causal laws. (2002, 61-62)

Not surprisingly, Kornblith believes a reliabilist account of knowledge coheres quite well with this theory of natural kinds.

One might worry that this methodological approach will lead to a theory that will substantially depart from our folk epistemic notions, and we will change the subject matter altogether. For once we look outward entirely toward the physical objects, who knows what sort of theory we will end up with. But Kornblith believes we won't be changing the subject because he anchors the epistemological investigation using a causal/historical account of reference. Since the referent doesn't change throughout the empirical investigation, since the referent is secured not by something internal to the subject but by external causal connections, one's epistemic beliefs can change quite a bit without changing the subject; the investigator will still be anchored to the subject of his investigation—knowledge or justified belief. Although Kornblith takes this to be a virtue of his theory, we will see how this could turn into a vice.

Kornblith's Unjustified Dismissal of Internalism¹¹

Kornblith admits that skepticism could be true, even global skepticism—the belief that we don't know anything whatsoever.

The phenomenon we call knowledge must have a certain degree of theoretical unity if reference is to be secured. Were we to discover that there is no more theoretical unity to the various items we call knowledge than there is to the set consisting of ships and shoes and sealing wax, then a presupposition of the introduction of the term would be undermined, and the view that there is no such thing as knowledge would be sustained. (2002, 23)

There is an intriguing and unique problem that Kornblith's view encounters once he admits the possibility of skepticism. This problem could weaken his naturalistic methodology.

Kornblith's view appears to imply the following:

1. If knowledge is not a natural kind, then the term 'knowledge' does not refer.
2. If the term 'knowledge' does not refer, it is an empty term.
3. Hence if knowledge is not a natural kind, then 'knowledge' is an empty term. That is, if knowledge is not a natural kind, then there is no knowledge.
4. Hence either knowledge is a natural kind or there is no knowledge.

Kornblith's naturalism commits him to the disjunction in 4. We believe 4 is false (and at least one of the premises) for the following reasons.

Kornblith's methodology must allow that 'knowledge' could fail to refer to a natural kind. Now, suppose we came to believe, after applying Kornblith's methods, that 'knowledge' does not pick out a natural kind and thus there is no such thing as knowledge. This would make a version of skepticism true.¹² Where would this leave us? Well, if we arrived at the belief there is no knowledge, we should then doubt whether we knew in the first place that either knowledge is a natural kind or nothing at all (and thus should be investigated using the natural kinds method). And if we have a reason to doubt that knowledge is a natural kind or nothing at all, then we have a reason to doubt whether skepticism is true in the first place, since knowledge could be something else (perhaps a non-natural kind). And if knowledge could be something else, then an internalist epistemic theory could be true. In this situation, coming to the conclusion that skepticism is true would warrant a fresh appraisal of how we went about searching for knowledge. We should question whether we were right to use the natural kinds approach. In other words, the search began by eliminating internalism as a viable contender for an analysis of knowledge or justified belief, but concluding that skepticism is true would not imply that our epistemic work is done. Unlike traditional skepticism, where at this point, we would have very little hope, Kornblith's kind of skepticism gives us a reason to start over and use a method that is more theory neutral—one that would take into account both naturalism and non-naturalism.

We can learn several things from the above. As we've just argued, Kornblith's methodology must allow for the possibility that internalism is true. But his methods cannot be used to further investigate internalism or non-naturalism. If, on the other hand, there is a philosophical method that both allows for the possibility of internalism and externalism and that method can be used to either investigate internalism or externalism, the use of that method is more justified. What is odd is that his methods make possible

our eventually arriving at a justified belief in internalism, but those very methods don't allow us to investigate internalism in any way.

So Kornblith's methodology *prima facie* rules out internalism (even if we conclude there is no knowledge). For him, either knowledge is a natural kind or there is no knowledge. Clearly, he won't discover internalism is true using the natural kinds approach. If he were to conclude that 'knowledge' does not refer, he would have no reason to go on to consider whether internalism is true. His naturalistic method would not allow for this. But actually, by countenancing the possibility that knowledge is not a natural kind but a gerrymandered kind, and given our previous arguments about what follows if we conclude there is no knowledge, he should allow that there is a legitimate possibility internalism is true. In other words, if there is a legitimate possibility that knowledge does not form a natural kind, and if this leads to the legitimate possibility that internalism is true, then he doesn't rightly restrict the outcomes to 4 (either knowledge is a natural kind or there is no knowledge).

We have already gestured at what we believe is the most penetrating objection to the naturalistic approach to philosophy. Earlier we worried that if we investigate knowledge empirically, we may find that our resulting empirical beliefs get further and further away from our ordinary intuitions about knowledge. We may find that knowledge is completely unlike anything we had previously believed. What then makes us think we're even speaking about the same subject matter? Since Kornblith believes we refer to objects not because of our beliefs about them or because of what is going on inside our heads but because of causal connections that we bear to them, once the referent is achieved via the right sort of causal connection, our beliefs can change and we won't be changing the subject—we won't lose track of the referent. If we work out the implications of the causal theory of reference when applied to 'knowledge', we get some unintuitive results.

Torin Alter and Russell Daw, in a little known article entitled "Free Acts and Robot Cats" (2001), raise the crux of the problem. The context of their discussion is free will. There they discuss Mark Heller's (1996) proposal to treat free will as a natural kind, and 'free will' as a natural kind term subject to the causal theory of reference. So, on this approach to philosophy, we learn the essential nature of free actions by empirical investigation, not conceptual analysis.

Alter and Daw give several objections, and although one in particular creates much trouble for Kornblith, permit us quickly to mention another of theirs. Imagine that through conceptual analysis we discover that our concept of justification is evidentialist and quite demanding on the human mind, requiring information processing tasks and the recognition of inferential connections many of us often fall short of. Suppose most

philosophers agree our concept is evidentialist. However, by using Kornblith's naturalistic approach, we discover justification forms no natural kind (ignoring the self-undermining nature of this claim for the moment). We later discover a group of Martians who satisfy our demanding concept of justification. Shouldn't we conclude that Martians have justified beliefs? And if so, then doesn't our concept of justification play an important role in epistemology? If a positive appraisal of their beliefs makes sense, even in the absence of discovering a natural kind, then the discovery of a belief's positive epistemic status can't depend (or at least, not entirely) on Kornblith's naturalistic methodology.

More technically, the natural kinds methodology implies the following (this is similar to the numbered argument above on pg. 6 but we're now interested in a different implication):

1. If T is a Putnamian-kind term, then if T refers it refers to a kind.
2. If T does not refer to a kind, then T fails to refer.
3. It is possible that knowledge does not form a kind.
4. Hence, it is possible that 'knowledge' fails to refer.
5. If 'knowledge' fails to refer, then there is no knowledge.

However, it is possible that "On Mars or in Heaven there are creatures [whose mental states fit our concept of knowledge]" (Alter and Daw 2001, 350). Nevertheless, such creatures would fail to know, since, according to the causal theory of reference and the assumption that knowledge does not form a kind, 'knowledge' fails to refer. There is no knowledge to be had by anyone (or anything).

The more significant problem is the possibility that the kind Kornblith may end up with has nothing to do with our concept whatsoever. If we ignore our concepts and focus entirely on the empirical realm, we could end up with an analysis completely incongruous with our ideas of what the thing is. For instance, consider the thought experiment proposed by Peter van Inwagen (in another context):

(M) When any human being is born, the Martians implant in his brain a tiny device...which contains a "program" for that person's entire life: whenever that person must make a decision, the device *causes* him to decide one way or the other according to the requirements of a table of instructions that were incorporated into the structure of the device before that person was conceived. (quoted in Alter and Daw 2001, 350)

To see what problems the natural kinds approach runs into, suppose that free acts form a natural kind, that 'free-action' is a natural kind term, and that all of our uses of the term

‘free action’ are causally connected to the kind of actions described in M above. Alter and Daw draw out the following menacing consequence for the naturalist¹³ “If ‘free action’ were a Putnamian-kind term, then the discovery that M is true would warrant the conclusion that free action is M-type behavior” (2001, 351). This result looks fairly implausible. Clearly their actions shouldn’t count as free. There is a paradigmatic case of unfree action. The authors go on: “Further, consider van Inwagen’s Martians, the creatures who implant the devices in the brains of human infants. Might their actions be free? If Heller’s proposal were true, then their actions would be free only if they were of the same kind as the human actions described in the M-scenario. We find that absurd as well” (2001, 351).

The same untoward consequence arises for Kornblith’s naturalism. If ‘knowledge’ is a Putnamian-kind term, then the nature of knowledge can only be discovered empirically. It is possible that in each application of the term ‘knowledge’, the alleged knowledge includes false beliefs, or lacks justification (or a similar epistemic property), or even lacks an intentional state altogether. The point is that knowledge could, given Kornblith’s methodology, turn out to be something completely different from anything we thought. Just as in van Inwagen’s scenario in which all of our uses of the term ‘free action’ actually pick out actions controlled by Martians, all of our uses of ‘knowledge’ may actually pick out mental states completely unlike anything we meant by the term. This looks implausible.

Traditional epistemology, on the other hand, does not face the same problem. There is not, at the beginning of their investigation into knowledge, the specter of possibilities that face the methodological naturalist. This is at least true in part because traditional epistemologists begin their analyses with different items. The methodological naturalist begins with empirical beliefs whereas the traditional epistemologist begins with items internal to one’s mind. Because of their access to their concepts, the traditional epistemologist can be pretty confident certain possibilities won’t obtain. To use a more extreme example, just think of all the common sense epistemologists who deny that local skepticism (skepticism with respect to a certain class of propositions) has any plausibility. They think they already know before even doing an analysis of knowledge that skepticism about the physical world has got to be false. If they think of themselves as analyzing a mind-dependent concept, then their view has some plausibility. They have better epistemic access to their own concepts than the essences of natural kinds, just as we have better epistemic access to our own concepts than to yours. But if they’re going after a mind-independent natural kind, their degree of confidence makes much less sense.

Kornblith perhaps isn’t sensitive to this objection because, in a sense, he lucks out. He believes he discovers an analysis of justification and knowledge already on offer—reliabilism. But before he begins, for all he knows, he could end up with a radically

different account of knowledge and justified belief. And given his methods, he will be beholden to that account, regardless of how absurd or anomalous the result is.

Zeroing in on the Right Properties

Even though some naturalists model the search for knowledge on natural kinds, the method for discovering philosophical objects is still different than the one used for garden variety natural kinds. We don't stumble over knowledge or see a piece of knowledge glinting in the sunlight and investigate the thing empirically. We could discover that rocks fall into a kind without any prior idea or notice of them. We could simply attend to them one day and recognize that several specimens share properties in common. Because of this difference between knowledge and other natural kinds, Kornblith must link what he does in some way to what knowledge truly is. Not anything can count as knowledge.

Earlier we noted that Kornblith thinks of intuitions as empirical judgments, no different in kind than the judgment that *x* is a rock. Several worries arise. First, once we've located the samples, how do we know which parts to focus on? Where should we look? Inside the brain? At the relation between the subject and the world? At the eyes? We know, assuming naturalism, that we should rely on sensory experience, but to discover the nature of knowledge, what precisely should we be experiencing? There are too many things in one's visual field to focus on, and we need some way of narrowing down the possible candidates so that we get knowledge, and knowledge alone, in our sights. The advice, "Just look for a natural kind," would provide little help, for there may be more than one natural kind in front of us.

At this point, a noteworthy problem in the philosophy of language appears—a version of the *qua* problem.¹⁴ A related epistemic problem crops up too, and we'll argue that Kornblith should accept that our concepts do much more to determine what we focus on and how we investigate than the naturalist admits. In other words, conceptual analysis will need to play a much larger role. We may need to rely on more general intuitions—intuitions that specify some of the characteristics of knowledge. This amendment would take us some way back to traditional epistemology.

According to (what we will call) the pure causal theory of reference (henceforth PCTR), singular terms and some general terms can refer even when the reference-grounder and reference-borrower don't have any correct descriptive content in mind. Because of this, PCTR quickly runs into a problem. For any naming ceremony, the referent is a member of numerous kinds (both natural and gerrymandered). For the term to pick out the species (in the case of natural kind terms) or the individual (in the case of proper names) and

nothing else, reference-grounding must include at least some correct descriptive content. If there isn't any, then the term could refer to any of the numerous kinds of which the referent is a member.

For instance, imagine that S comes into contact with a species and baptizes the members as 'tigers'. PCTR tells us that S's initial baptism refers to all and only tigers. But when S confronted tigers, S also confronted cats, mammals, animals, time slices of tigers, individual tigers, etc. How could S's initial baptism with 'tiger' be a baptism of all and only tigers and not cats, or mammals, or animals? The fact that members of natural kinds belong to an infinite number of kinds—both natural and gerrymandered—suggests that PCTR is incomplete. Nothing in PCTR rules out that with S's initial baptism, S baptized a different kind or even a disjunction of different kinds.

This is the qua-problem. We need an account of S's initial baptism that explains how S managed to refer to tigers qua tigers and not tigers qua cat, mammal, or something else. Likewise we need an account from Kornblith that explains how we manage to refer to knowledge qua knowledge and not knowledge qua mental state, knowledge qua belief state, or what have you.

We can draw out a similar worry by exploring the dependence of the natural kind methodology on empirical intuitions and empirical samples. As we already asked, when we get enough samples in front of us, how do we learn where to look for knowledge? The samples will often share too many properties in common.

Nevertheless, we have a pretty good idea where to look. We're not stumbling around scratching our heads wondering what in the world this thing 'knowledge' could be. Animal? Vegetable? Mineral? Probably, two things occur here to help us out. First our concepts generate intuitions that lead us to concentrate on certain properties and second, these intuitions tend to be more general in form.¹⁵ As a rule, philosophers rely on general intuitions more than they think. Even traditional philosophers don't appear to rely solely on particular case intuitions to get the cases out into the open. If the various cases of justified belief have quite a few properties in common, we have an intuition about which properties to isolate. Furthermore, even when we find ourselves sharing many of the same intuitions, we may still find our theories at loggerheads. This is probably because of our intuitions about the more general epistemic and moral truths.

So to figure out where to look when we've collected samples of knowledge, our more general intuitions resulting from concept possession must play a larger role. If we're right and the natural kinds approach suffers from something like the qua problem, then we need to rely on our concepts more than Kornblith initially suggests. His extra-mentalist view makes it possible that concepts misrepresent. Concepts, for him, can be more or less

mistaken.¹⁶ Now even if our concepts can be mistaken, we can still do a pretty good job picking out samples of something. Even those who have the most arcane views about water don't have much of a problem using water to quench their thirst. Consequently, even if our epistemic concepts have been influenced by many of the right cases, there isn't much reason to think our concepts do a pretty good job representing the essence of knowledge qua natural kind. Thus naturalists should be wary of trusting more general intuitions unless they have reasons to think our concepts pretty closely correspond to the natural kind. Here then we have an epistemic version of the qua problem. Instead of focusing on how reference gets underway, now our focus is on how the epistemic investigation gets underway.

Perhaps this explains why in chapter 2 of his book, Kornblith seems to shift gears. Given his arguments in chapter 1, the reader expects to find obvious case after obvious case of knowledge to think about empirically. He doesn't get samples of human knowledge before him and try to discover an underlying theoretical unity. The discussion shifts to animals. He tries to explain animal behavior, ultimately arguing that the best explanation for certain kinds of animal behavior is that they have knowledge. He moves to using inference to the best explanation, which we often use when we can't see the thing we're trying to investigate, but we wish to argue is there.

So here is a quick summary of the argument: Natural kind theories of knowledge quickly invite the epistemic qua problem. So we must find which properties to focus on. But then we'll probably need to rely on our more general intuitions (intuitions influenced by our concepts). But, if knowledge is a natural kind, we don't have much justification for thinking our epistemic concepts will closely correspond to the essence of knowledge. Thus we won't have much justification for thinking we're singling out the right properties.^{17, 18}

This problem becomes more magnified when we call into question the reliability of our (empirical) intuitions about particular cases. Note the contrast between gathering samples of water or heat and gathering samples of knowledge. Our intuitions tend to be far less uniform in the epistemic realm.¹⁹ This is probably because we identify water and heat through observable properties. They have a certain group of properties that uniformly influence our classificatory judgments.

Naturalists like Kornblith, Heller, and Richard Boyd²⁰ often draw parallels between the types of methods they endorse in the search for knowledge, free will, and moral goodness and the search for the nature of such things as rocks, heat, and water. But the analogy is unhelpful, for we know how to pick out rocks, heat, and water. We can see and touch such things. We know how to isolate them. We felt heat, we stumbled over rocks, we drank and swam in water. But with goodness, free will, and knowledge, simply relying on

our common sense judgments about particular cases hardly seems reliable (assuming the natural kinds view). When we say, “Look, there is a case of knowledge,” you can’t be quite sure what we’re trying to draw your attention to. But in the case of water, at least we can splash some on you.

If we’re right, this, in part, explains why there is so much disagreement in philosophy. If knowledge had observable properties that we could clue in on, there would probably be much less disagreement among our particular case intuitions. From the natural kinds point of view however, it’s not terribly likely that our intuitions about knowledge are somewhat reliable. Since knowledge doesn’t have observable properties, there doesn’t seem to be anything that would make likely our being reliable in our epistemic intuitions about cases. If on the other hand you believe that those intuitions proceed from concepts, concepts that determine the structure of philosophical phenomena, there is much more reason to trust the resulting intuitions.

Obviously, we need our intuitions to be reliable. But the reason is initially more pressing for a natural kinds metaphilosophy. First off, if our intuitions are only somewhat trustworthy, it’s more likely that we will start out with the wrong cases. We might then search for a theoretical unity among them and think we’ve found knowledge. This is a worry for any philosopher. But the traditional philosopher has a check against this possibility not open to Kornblith. As Kornblith says:

The judgments of rock collectors at early stages of investigation are substantially inferior, epistemically speaking, to those at later stages, when theoretical understanding is further advanced. We should not say that initial intuitive judgments are of no evidential value, for were this the case progress in theory would be impossible...but, that said, it must also be acknowledged that judgment guided by accurate background theory is far superior to the intuitions of the naïve. Intuition must be taken seriously in the absence of substantial theoretical understanding, but once such theoretical understanding begins to take shape, prior intuitive judgments carry little weight unless they have been endorsed by the progress of theory. (Kornblith 2002, 14)²¹

One of the difficulties is that Kornblith wants to get the empirical investigation under way, to start leaving intuitions behind and isolate the relevant properties. Which means we should begin to rely more on theory mediated judgment. So intuitions that occur later on that clash with the theory—a theory based on empirical investigation—or that suggest the initial examples are wrong, should have less epistemic weight. But this makes starting off with the right samples even more important, for there isn’t the same check in place as there is for the traditional philosopher. The traditional philosopher can always appeal to a

stronger intuition later in the philosophical process which can act as a check on his earlier intuitive judgments. Traditional foundationalists will often make use of this. They might begin by having intuitions that correspond to those of the common sense philosopher, but then later they might come across a principle, like Richard Fumerton's principle of inferential justification, that seems more intuitive to them than the initial examples. Armchair philosophers usually worry that there will be some intuition later in the theoretical process that will undermine their theory or cast doubt on their most cherished particular cases. These intuitions don't proceed from their theory since they speak against the theory. Even when we get well-developed theories, our intuitions aren't well enough behaved that we know later intuitions will always be trumped by theory.

On our view, initial intuitive judgments don't possess as much weight as the later intuitive judgments that occur after we have taken in quite a bit of the relevant information, know more of the relevant scenarios, the moves open to us, and so forth. When one is further along in the process of adjusting theory to intuitions and intuitions to theory, one's intuitions have more evidential weight.²² These later intuitions act as an important check on our earlier intuitions giving us reason to think our earlier intuitions furnished us with the right cases. This isn't available for the natural kinds approach. Initial intuitive judgments should give way to judgments influenced by theory. If one has an intuition much later in the process that suggests the examples we began with were wrong (or at least we don't know they're right), this intuition will compete against the developing theory. The theory is drawn from empirical evidence, a scientific investigation into a natural kind. Who knows where this later intuition came from? So theory should trump intuition, according to the natural kinds philosopher.

Michael Huemer raises a similar problem (2008, 84-87). Huemer argues that given the way these natural kind philosophers talk, we expect that goodness or knowledge is supposed to be analogous to heat or water, since the philosopher hopes to explain the underlying nature of these philosophical objects using natural properties, just as scientific theories explain the underlying nature of heat and water using natural properties. "This suggests," according to Huemer, "that moral concepts should be observational, just as 'heat', 'water', and 'sound' are. The reason why we believe, for example, that $\text{water}=\text{H}_2\text{O}$ is, roughly, that...we have independent (that is, pre-theoretical), direct awareness of the presence of water" (2008, 85). So applying the naturalist strategy, we need moral goodness and knowledge to look like something, feel like something, or have some observable characteristics such that we can isolate them as we do water and heat. Again, one doesn't find some knowledge glinting in the sunlight. This is why the attempt of naturalists to model their investigation on the search for rocks or gold sits awkwardly in the mind.

We suspect then that our epistemic beliefs, on the naturalists' view, dictate far more than they ought; they determine a certain answer to the question "What is knowledge?" Instead of letting the empirical data do the talking, our concepts and our beliefs about what knowledge is limit the field of inquiry allowing us to circumscribe the number of objects under investigation and pointing us to places to focus even after we have picked out our samples of knowledge.

Kornblith on Knowledge

As we noted earlier, rather than rely on standard examples of knowledge and then proceed to discuss the empirical information that has been gleaned by scientific investigation thus far, Kornblith makes a peculiar move when he proceeds to chapter 2 of his book. He begins by analyzing animal knowledge.

There is a large literature on animal cognition, and workers in this field typically speak of animals knowing a great many things. They see animal knowledge as a legitimate object of study, a phenomenon with a good deal of theoretical integrity to it. Knowledge, as it is portrayed in this literature, does causal and explanatory work...if cognitive ethologists are even roughly right, then talk of animal knowledge is not a mere *façon de parler*; rather, there really is such a thing as animal knowledge. Knowledge constitutes a legitimate scientific category. (2002, 28-29)²³

As one can see, he thinks that we can learn quite a bit from recent studies of the mental lives of animals. But we're not sure what entitles him (even assuming the truth of his view) to begin with animal knowledge since attributions of propositional knowledge to animals is not something most epistemologists agree on, at least to the sorts of animals he considers (like ants for instance). We doubt this has curried enough favor among epistemologists to warrant the study of certain examples of alleged animal knowledge. And even if there is some agreement about these cases, there is still too much disagreement about animal knowledge to provide Kornblith with enough canonical examples. Furthermore, he does not use an inductive approach but relies on argument to the best explanation. The behavior of animals is best explained by attributing knowledge to them. But shouldn't this come only after we have figured out what knowledge is? How does he know knowledge is present? He can't say, "Let's call this thing 'knowledge'." These are issues we will explore as we observe how Kornblith reaches his reliabilist account of knowledge.

Kornblith first argues that the behavior of animals exhibits a complexity rich enough to demand that we attribute to them intentional states like beliefs. We can't distinguish, say,

flight and play without attributing to them intentional states. He goes on to argue that the only way to find what is common to instances of animal behavior is by attributing intentional states to animals. There is no commonality at the level of bodily motion. The behavior of one human distracting another may have nothing in common at the level of bodily motion with other instances of humans distracting others. Only mental states can unify the activity. And the same goes for animals. The difference between flight, running, or playing must refer to the reasons for the behavior since there may be no difference at the level of bodily motion.

For the next step Kornblith introduces representation of the environment. He claims that an animal, “If it is to satisfy its biologically given needs, it will need to recognize certain features of its environment and the evolutionary process must thereby assure that an animal has the cognitive capacities that allow it to deal effectively with that environment. What this requires is the ability to represent information” (2002, 37). He goes on: “Once we recognize the existence of internally represented animal needs together with representations of features of the environment, we have the beginnings of a belief-desire psychology. The ravens distract the hawk because they are hungry; they want to steal the hawk’s egg; they believe that by attempting to take the squirrel away from the hawk, they will thereby be able to take the egg” (2002, 38).

Kornblith then approaches his conclusion by considering species in general:

[W]e are interested in an explanation of how it is that members of the species are endowed with a cognitive capacity that allows them successfully to negotiate their environment. It is the focus on this adaptation of these cognitive capacities to the environment that forces us to explain the possibility of successful behavior, and it is the explanation of successful behavior that requires the notion of knowledge rather than mere belief. (2002, 57)

Of course, successful behavior requires reliable belief production. One must be forming beliefs in a reliable way to successfully navigate their environment and survive:

If we are to explain why it is that plovers are able to protect their nests, we must appeal to a capacity to recognize features of the environment, and thus the true beliefs that particular plovers acquire will be the product of a stable capacity for the production of true belief. The resulting true beliefs are not merely accidentally true; they are produced by a cognitive capacity that is attuned to its environment. In a word, the beliefs are reliably produced. The concept of knowledge which is of interest here thus requires reliably produced true belief. (2002, 58)

One natural question to ask upon reading Kornblith's argument is how knowledge suddenly appeared on the scene. This leads to a related point: cognitive ethologists may not know all the relevant epistemic concepts available to them for explaining animal behavior. Philosophers, for one, may want to make other more subtle distinctions perhaps between knowledge *de re* and knowledge *de dicto* or knowing how and knowing that.

At the end, once he finds reliable belief-forming mechanisms, he says this is knowledge. But he has left the door open for the internalists, for they will deny we have found knowledge. They will wonder how all the sudden knowledge entered the picture. By using inference to the best explanation and waiting until further along in the process to make a knowledge pronouncement, Kornblith appears to let his concept play a substantial role in what counts as knowledge. For remember what he told us originally in chapter 1: we pick out items of knowledge and then we discover their nature. But here we discover the nature of something and we end up calling it 'knowledge'. But why call it knowledge rather than mere reliable belief? When using inference to the best explanation, we have options open to us as to what to call this thing which helps animals navigate their environment, and other concepts will do just as well.

Consequently, he is not allowing the empirical evidence to inform his beliefs about knowledge but rather is taking his beliefs about knowledge and imposing them on the empirical evidence.

Bibliography

Alter, Torin and Russell Daw. 2001. Free Acts and Robot Cats. *Philosophical Studies* 102, no.3 (February): 345-357.

BonJour, Laurence. 1999. *In Defense of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge: Cambridge University Press.

Boyd, Richard. 1997. How to Be a Moral Realist. In *Moral Discourse and Practice*, edited by Stephen Darwall, Allan Gibbard, and Peter Railton. New York: Oxford University Press.

Goldman, Alvin. 1988. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

----- . 2005. Kornblith's Naturalistic Epistemology. *Philosophy and Phenomenological Research* LXXI, No. 2 (September): 403 – 410.

Heller, Mark. 1996. The Mad Science Meets the Robot Cats: Compatibilism, Kinds, and Counterexamples. *Philosophy and Phenomenological Research* 56, no2 (June): 333-337.

Huemer, Michael. 2008. *Ethical Intuitionism*. New York: Palgrave MacMillan Ltd.

Jackson, Frank. 1998. *From Metaphysics to Conceptual Analysis*. Oxford: Oxford University Press.

Kornblith, Hilary. 1998. The Role of Intuition in Philosophical Inquiry: An Account with No Unnatural Ingredients. In *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, edited by Michael R. DePaul and William Ramsey. Lanham, MD: Rowman and Littlefield Publishers, Inc.

-----, 2002. *Knowledge and It's Place in Nature*. Oxford: Oxford University Press.

-----, 2007. Naturalism and Intuitions. *Grazer Philosophische Studien*, 74: 27 – 49

Sosa, Ernest. (2005). A Defense of the Use of Intuitions in Philosophy. In *Stich and His Critics*, edited by Michael Bishop and Dominic Murphy. New York: Blackwell Publishers.

¹ We will not be defending our own view here. This is a critique of a metaphilosophy and not a defense of any specific alternative metaphilosophy. We will, at times however, contrast Kornblith's view with the traditional approach. Thanks to an anonymous referee for pointing out our need to do this.

² An alternative is furnished by the Aristotelian/ Thomist. Words are signs of conceptions which, in turn, are signs of beings (either accidental or substantial beings, or transcendentals). According to this view then, philosophy is about the extra-mental world, not just our representations of that world. But we get to the world through the signs for those things—our conceptions. We mention this because this is the metaphilosophical approach we lean toward.

³ This is too strong for our tastes since we believe there must be good things (a Thomistic view), and in any world with beings that have cognitive capacities like our own, there must be knowledge. This doesn't make much of a difference though since our view is a long way from Kornblith's, even further than mentalist views that admit the possibility of a world with no good things.

⁴ As we'll soon learn, Kornblith's view admits the possibility that 'freedom' doesn't have any conditions for its correct application; freedom could be nothing at all.

⁵ Or consider those atheists who are metaethical theological voluntarists. They believe we can learn that legitimate moral rules depend on God even in a world where God does not exist.

⁶ See his (2007) pg. 36.

⁷ “Here I simply take for granted a causal or historical account of reference of natural-kind terms. While the details of such a theory remain to be established, the general outline is, I believe, perfectly clear in the foundational work of Kripke and Putnam,” (Kornblith 2002, 12 fn. 18).

⁸ See Goldman (1988). In his later work he still emphasizes the importance of conceptual or meaning analysis but grants more of a role to the sciences, especially psychology, in this metaepistemological stage.

⁹ The closure principle is just one of the examples in epistemology. Moral pluralism is probably another case of an ethical theory that stresses the importance of moral principle intuitions. Keep in mind that the conceptual analyst should probably accept this idea. If concepts ‘in the mind’ can spark intuitions about particular cases, they probably can engender general intuitions about the principles themselves (principles that tend to be encoded in the concepts).

¹⁰ In his (2007) he further explains his view of the relation between theory and theory influenced intuitions. His discussion tackles some interesting problems.

¹¹ A number of critics—in particular Frank Jackson, Alvin Goldman, and Ernest Sosa—have raised concerns similar to the some of the objections raised in this paper. But we attempt to flesh out and improve upon those worries raised in the earlier literature. Thanks to an anonymous referee for drawing this to our attention. For a sampling of the literature, see Jackson (1998), Goldman (2007) (among others) and Sosa (2005).

¹² This is a different kind of skepticism than traditional skepticism since knowledge turns out to be nothing at all. If we discover knowledge is a gerrymandered kind, we aren’t talking about anything when we try to refer to knowledge. Usually epistemologists recognize that if skepticism is true, knowledge could still have conditions for its instantiation. But when we say, “there is no knowledge” on Kornblith’s view, we don’t simply have a reason to doubt that knowledge is instantiated or had by anyone. We have a reason to believe knowledge is nothing at all rather than thinking we just fail to satisfy the conditions for knowledge. We don’t discover that knowledge requires reliabilism and then later learn we don’t use reliable processes (though how we could learn this is something of a puzzle).

¹³ Unless we say otherwise, we’re using the terms ‘naturalism’ or ‘naturalists’ to refer to those who endorse the natural kinds approach to philosophical investigation.

¹⁴ Causal theories of reference tend to founder on this problem which leads causal theorists to supplement their theory with descriptive constraints in order for reference to be successful.

¹⁵ Particular intuitions tend to take the form that some scenario is or is not a case of knowledge. General intuitions tend to have more general principles in their content. They may even be about the conditions required for the instantiation of some concept.

¹⁶ Once again, see Kornblith’s (2007) for this characterization of concepts.

¹⁷ “The subject matter of the rock collector’s investigation is the natural kind, whatever it may be, which (most of) the samples picked out are members of; but the investigator need not be in a position to

characterize the essential features of that kind. The investigator's concept of that kind, therefore, because it may be quite incomplete or inaccurate, need not itself do very much of the work of defining the subject matter under study" (Kornblith 2002, 11).

¹⁸ We're not arguing that we must know what the nature of something is before we can pick out examples. We think both particular and general intuitions play a role in the reflective equilibrium process. Our point is that Kornblith must appeal to these more general intuitions and this creates some difficulties, since concepts can misrepresent kinds and we don't have much justification for thinking that appealing to the more general intuitions delivered by our concepts is a trustworthy process. When concepts can misrepresent and when we're trying to represent a natural kind, we have less justification for trusting the content of our concepts than the traditional mentalist does.

¹⁹ Some might argue that we still have uniform intuitions on enough cases to furnish us with enough samples. But we're skeptical of this point. Traditional foundationalists tend not to share our intuitions about particular cases. What might be an obvious case of knowledge to us is unlikely to be to them. What then do we do with their intuitive responses? Ignore them? These aren't your average person making epistemic judgments. These are people Kornblith would describe as having thought long and hard about epistemic issues. How then can we say that we have collected a lot of samples the experts would agree on? It's not as though philosophers have a different particular case intuition here or there. Rather, they run into particular cases they don't agree on rather quickly. Nor do the experts agree on a certain set of particular case intuitions. There are epistemologists throughout philosophy's history and even now who would disagree with items in that set.

²⁰ See his "How to Be a Moral Realist" (1997).

²¹ Kornblith also says that "While it is certainly true that judgments driven by bad theories are not to be taken seriously, the solution is not to try to return to some pure state of theory-independent judgment, before the fall, as it were; rather the solution is to get a better theory. Intuition in the absence of theory does not count for nothing, especially if no credible theory is available. But this is not to award high marks to intuitive judgment before the arrival of a successful theory, let alone after, when the initially low value of such judgment drops still lower" (2002, 14).

²² Earlier intuitive judgments are more like walking in half-way through a movie and being asked to explain the plot.