

EPISTEMOLOGY AND EMPIRICAL INVESTIGATION¹

Kristoffer AHLSTROM
UMass Amherst

Summary

Recently, Hilary Kornblith has argued that epistemological investigation is substantially empirical. In the present paper, I will first show that his claim is not contingent upon the further and, admittedly, controversial assumption that all objects of epistemological investigation are natural kinds. Then, I will argue that, contrary to what Kornblith seems to assume, this methodological contention does not imply that there is no need for attending to our epistemic concepts in epistemology. Understanding the make-up of our concepts and, in particular, the purposes they fill, is necessary for a proper acknowledgement of epistemology's role in conceptual improvement.

1. *Introduction*

In his book *Knowledge and Its Place in Nature* (2002), Hilary Kornblith makes an intriguing case for the re-conceptualization of epistemological analysis from a largely non-empirical to a substantially empirical investigation, arguing that knowledge—one of the main targets of epistemological investigation—is a natural kind, open to straightforward empirical scrutiny. Assuming that knowledge is not unique in this respect, which is an assumption that Kornblith, indeed, seems to make, we may generate the following argument:

The Argument.

(A) All objects of epistemological investigation are natural kinds.

1. Earlier versions of this paper were presented at the University of Massachusetts Amherst as well as at the 2007 meeting of the Danish Epistemology Network in Copenhagen, May 2007. I would like to thank both audiences for valuable discussions and comments, and am particularly indebted to Hilary Kornblith, Klemens Kappel, Joseph Levine, Åsa Wikforss, Helge Malmgren, Radha Vij, Alex Sarch, and Kelly Trogon.

- (B) If (A), epistemological investigation is substantially empirical.
- (C) Hence, epistemological investigation is substantially empirical (A, B, *MP*).
- (D) If (C), a thorough understanding of our epistemic concepts, over and above the phenomena that they pick out, is irrelevant to epistemological investigation.
- (E) Hence, a thorough understanding of our epistemic concepts, over and above the phenomena that they pick out, is irrelevant to epistemological investigation (C, D, *MP*).

Kornblith has never explicitly stated this argument. Still, I take it that it provides one of the most reasonable rationales for his more general claims about the implications of his results concerning knowledge to epistemological analysis at large.² The plausibility of this interpretive claim should become more obvious as we go along.

That being said, I will, in the following, scrutinize, qualify, and criticize *the Argument* in two steps. More specifically, §§ 2 through 4 will serve to contest (A) but defend (C), by showing that the latter premise is plausible even given that all objects of epistemological investigations are artificial (or “socially constructed”) rather than natural kinds. However, §§ 5 through 7 will show that (E), nevertheless, does not follow from (C), since (D) is false and the claim that epistemological investigation is substantially empirical hence, does *not* imply that an understanding of our epistemic concepts is irrelevant to epistemology.

2. *On the implausibility of premise (A)*

It should be beyond doubt that *the Argument* is valid. Indeed, it consists in two *modus ponens* arguments, where the conclusion of the first, i.e., (C), makes up the first premise of the second. However, I would like to contest its soundness. For one thing, it hinges on (A), i.e., the controversial assumption that all objects of epistemological investigation are natural kinds. As already mentioned, Kornblith (2002) has, indeed, argued that knowledge, as it is being studied by cognitive ethologists (cognitive ethology being the study of animal cognition), is a natural kind. However, the crucial question here is whether this claim may be

2. See, e.g., Kornblith (2007, 2006).

generalized to other objects of epistemological study, so as to render (A) plausible.

Take epistemic justification, for example. What are the prospects for extending Kornblith's case for knowledge to justification? Unfortunately, unlike KNOWLEDGE, JUSTIFICATION is not an entrenched concept in cognitive ethology. Hence, it is questionable whether Kornblith's case for knowledge can be extended to justification in any straightforward way. In fact, it is hard to see exactly how JUSTIFICATION, together with such related concepts as EVIDENCE, UNDERSTANDING, and RATIONALITY, at all *could* correspond to natural rather than artifactual (or "socially constructed") kinds, the latter of which do not comprise, say, homeostatically structured conglomerates of properties independent of human understanding, but rather a grid whose structure reflects nothing but human intentions.³ Still, as has been noted by Alvin Goldman and Joel Pust (1998), the lack of natural kind status hardly places the topic of justification (or that of evidence, understanding, and rationality) outside the scope of epistemological analysis.

So, on pain of radically restricting the scope of epistemological analysis (an option that I will not consider), the defender of *The Argument* has to face up to the following problem:

Problem 1.

Unless (A) holds, there is little reason to believe that the applicability of epistemological analysis stretches beyond the analysis of one particular object of epistemological investigation, namely knowledge. And even this particular application is contingent upon the admittedly controversial claim that knowledge, in fact, is a natural kind.

The following two sections discuss two solutions to *Problem 1*, both of which amount to the claim that there is a case to be made for extending the conception of epistemological analysis as substantially empirical to the analysis of artifactual kinds.

3. For an account of natural kinds as homeostatically related properties, see Kornblith (1993).

3. *A first attempt to save (C): Content externalism about artifactual kind terms*

The first solution starts out with the observation that it might be plausibly argued—and, indeed, has been argued by Putnam (1975) and, more recently, by Kornblith (forthcoming)—that content externalism provides the correct semantics not only for natural kind terms but also for artifactual kind terms. Rather than directly contesting this line of argument, the second solution (which is the one I will favor) concludes that, as it turns out, the plausibility of extending the claim about empirical investigation to artifactual kinds is largely independent of which semantic theory one accepts for the latter. Before evaluating any of these solutions, however, we need to say something about what constitutes artifactual kinds and, in particular, what distinguishes them from natural kinds.

To a first approximation, we may characterize artifactual kinds negatively as *not* comprising homeostatically clustered properties. However, even disregarding the fact that this characterization is hardly informative, it does not even uniquely pick out artifactual kinds, unless natural and artifactual kinds exhaust the realm of kinds (which they do not). So, by way of a positive characterization, we may say that artifactual kinds are somehow dependent on human intentions. However, this formulation is not only vague but also potentially misleading if not further qualified. Take polyethylene or amphetamine, for example. Since they are synthetic substances, it is plausible to assume that neither of them would be around if it were not for certain human intentions, pertaining to the need for a light, flexible, yet tough material or a substance to fight fatigue and increase alertness among servicemen.

Still, this only serves to show that the existence of some (instances of) synthetic substances are *causally* dependent on certain human intentions. It does not show, however, that the kinds to which those substances correspond are *ontologically* dependent on human intentions. That is, it does not show that the identity conditions for polyethylene or amphetamine—i.e., the conditions specifying what *makes* something an instance of polyethylene or amphetamine—are, in any interesting sense, intertwined with human intentions. Indeed, an acknowledgement of this very fact is implicit in what we take to be the best explanation of why polyethylenes and amphetamines fit into reliable inductive generalizations better than any random motley of properties. This explanation assumes that polyethylenes and amphetamines are endowed with a shared, underlying nature (i.e.,

C₂H₄ and C₉H₁₃N, respectively), and that this, furthermore, accounts for the fact that some inductions involving the respective substances are successful (e.g., from “this is amphetamine” to “this will increase stamina but decrease appetite if ingested”) while others are not (e.g., “this is made of polyethylene” to “this is blue”). Hence, they may plausibly be considered natural kinds.

Not so for, say, pens—a clear example of an artifactual kind. There is no need to assume that all pens share an underlying nature to explain why certain inductions involving pens are successful (e.g., from “this is a pen” to “this can be used to write with”) while others are not (e.g., from “this is a pen” to “this is warm”). The reason is that instances of artifactual kinds owe their kind membership exclusively to the fact that they fulfill certain purposes. More specifically, I suggest that

Identity Conditions for Artifactual Kinds.

the identity conditions of artifactual kinds are given by sets of human intentions, pertaining to the fulfillment of certain purposes.⁴

Clearly, this is not to say that artifactual kinds *consist of* sets of human intentions and purposes, but that what determines whether or not something is an instance of a particular artifactual kind pertains to whether that something can fulfill certain purposes and, thereby, answer to a specific set of human intentions.⁵ Thus, a pen is a pen (roughly) by virtue of fulfilling the purpose of drawing and writing and, thereby, answering to certain human intentions regarding creative outlet and communication, just like a key is a key (roughly) by virtue of serving the purpose of locking and unlocking doors, lockers, etc., and, thereby, answering to a set of human intentions regarding controlled access to certain spaces.⁶

Let us now consider Kornblith’s (forthcoming) claim that the semantic mechanisms of reference for artifactual terms are insensitive to these ontological differences between natural and artifactual kinds. Take an SUV, for example—clearly, an example of an artifactual kind. Unlike the case of

4. See Thomasson (2003) for a more detailed suggestion along these lines.

5. It might be argued that it, for some artifactual kinds, is not sufficient for membership that something merely can fulfill certain purposes and, thereby, answer to certain human intentions but that it has to have come about as the result of an intention to fulfill those purposes. See Thomasson (2003, 594) for a discussion.

6. I am not suggesting that such sets of intentions can be summed up in anything like a clear-cut conjunction of properties. This picture—just like actual categorizations of artifacts—is fully compatible with conceptual fuzziness and in-between cases.

water and polyethylene, there is no reason to assume that SUVs share an underlying nature, since an explanation of why we categorize the world and successfully reason in terms of SUVs and non-SUV type vehicles does not need to go beyond factors pertaining to certain (potentially superficial) properties regarding form (e.g., relative size) and function (e.g., performance), answering to certain human intentions concerning traveling and transportation. In fact, I am, personally, not sure what makes something an SUV and, in particular, not what distinguishes it (if anything at all) from a jeep, van or any other fairly big motor vehicle with four wheels. Regardless of whether I, thereby, just happen to be exceptionally uninformed as for motor vehicles, however, I take it to be fairly uncontroversial that I, nevertheless, just succeeded in referring to SUVs. How can that be?

To a first approximation, it may be due to the dual fact that (a) there are people in my linguistic community that *do* know what makes something an SUV and (b) my successful reference to SUVs is parasitic upon their knowledge and ability to discriminate SUVs from non-SUV type vehicles. But are these conditions *necessary* for successful reference? Is it, in particular, necessary that there is at least one member of my linguistic community that knows what, thereby, constitutes SUVs? Remember that what makes something an SUV pertains to a set of human intentions and purposes—not anything like an underlying nature, shared by all SUVs. So, is successful reference contingent upon there being at least one member of my linguistic community that knows what this set is, i.e., to which intentions SUVs need to answer and what purposes they need to fulfill? Consider the following two responses, corresponding to two variants of content externalism:

Social Externalism about Artifactual Kind Terms.

Successful reference to artifactual kinds (only) requires that there is at least one member of the relevant linguistic community (i.e., an “expert”) that can correctly delineate the set of human intentions and purposes that provides the identity conditions for the kind in question. Subsequent instances of successful reference to this kind are then parasitic upon the discriminatory competence of this member.⁷

Causal Externalism about Artifactual Kind Terms.

Successful reference to artifactual kinds (only) requires that a sample of the kind has been picked out in an initial act of baptism through an

7. See Burge (1986).

ostensive definition, fixing the set of human intentions and purposes that provides the identity conditions for the kind in question and establishing a socially sustained chain of reference upon which subsequent instances of successful reference to whatever bears a certain equivalence relation (most plausibly spelled out in terms of certain potentially superficial properties regarding form and function) to the ostended sample are parasitic.⁸

Clearly, neither formulation is intended to constitute a full-fledged theory. If anything, they both give rise to further questions. Take *Causal Externalism*, for example. For one thing, it is somewhat puzzling how the mere causal relation involved in an act of baptism could fix *one* set of human intentions and, consequently, pick out a *single* kind. For example, take a wooden, box-like, and fairly heavy object of 35 by 35 inches that we may name *b*. Picking out *b* as a sample for an artifactual kind, what determines the relevant set of intentions, given that *b* can be used to sit on (i.e., as a chair), to sit by (i.e., as a table), to stop doors from closing (i.e., as a door-stop), if pushed down a set of stairs, as a device stopping people from ascending (and possibly hurting them quite badly in the process of doing so), and probably a whole host of other things?

One plausible suggestion is that there is something about the mental state of the baptizer that determines the relevant set of intentions, namely that the mental state (at the very least) *instantiates* that very set of intentions. If so, however, it becomes harder to distinguish *Causal* from *Social Externalism*. For example, is successful reference contingent upon (a) the act of baptism and the resulting chain of reference, or (b) the fact that there is a baptizer, carrying the heaviest burden in the division of linguistic labor due to her insight into the relevant set of intentions (granted introspective access, of course)? It is not so clear to me which one is the case here.

However, the relevant question for our purposes is how to *analyze* artifactual kinds and, in particular, whether any of the above considerations render the claim that epistemology is substantially empirical implausible. They do not. The idea that epistemology is substantially empirical can be plausibly extended to artifactual kinds, regardless of whether *Social* or *Causal Externalism* holds and for the following reason: Both *Social* and *Causal Externalism* are fully compatible with successful reference despite substantial ignorance regarding the actual properties of the entities or

8. See Kornblith (forthcoming).

phenomena picked out. Granted, *Social Externalism* implies that there is at least one member of the relevant linguistic community that has insight into the identity conditions of the kind in question. So, in the general case, the baptizer, clearly, knows *something* about the artifactual kind she is baptizing, such as that it can be used for certain purposes and, thereby, answer to certain human intentions. However, her knowledge of many of the properties that constitute the artifactual kind in question may be ever so limited—or better said: there is nothing in her role as a *baptizer* that hinders her knowledge from being limited thus. That is, even if we reject *Causal* in favor of *Social Externalism*, the *epistemically most privileged user* of an artifactual kind term, i.e., the baptizer herself, may have an ever so limited insight into the properties that do or may make up instances of the kind in question.

As the reader surely suspects, it is exactly in this gap between successful reference and an insight into the properties of the referent picked out that our first solution to *Problem 1* gets its foothold, since such a gap makes possible scenarios in which (a) a majority of speakers either are largely *ignorant* of or have a highly *inaccurate* conception of many of the properties that make up the kind that they are successfully referring to and (b) even the most epistemically privileged speaker (i.e., the baptizer) may have an ever so limited insight into the multitude of properties that may not in any straightforward way be inferred from the intended purpose of the kind in question. Given such a gap, I take it that it would reasonably follow that epistemology is substantially empirical even given that most (if not all) objects of epistemological investigation are artifactual kinds.

4. *A second attempt to save (C): Epistemology and conceptual refinement*

The problem with this solution, however, is that it is far from controversial that anything like *Social* or *Causal Externalism* provides the correct semantics for artifactual kind terms. While remaining essentially neutral on this issue, and in an attempt to develop a somewhat more dialectically sensible rationale for (C), I will now argue that, *even if* a strong form of internalism turned out to provide the correct semantics for artifactual kind terms, this would in no way undermine the claim that epistemology is substantially empirical. The argument will also indicate that it was not externalism that did the job in the above solution after all.

So, consider the following:

Strong Internalism about Artifactual Kind Terms.

Successful reference to artifactual kinds requires that the speaker can correctly delineate the set of human intentions and purposes that provides the identity conditions for the kind in question.

Strong Internalism makes up the other extreme of the semantic spectrum; it is not enough that there is an appropriate causal chain of reference, nor that *someone* in the linguistic community can delineate the set of intentions in question—the speaker must *herself* be able to make such a delineation for her to successfully refer. Perhaps this is a more plausible thesis about the semantics of artifactual kinds, or perhaps it is not. What is important to note for our purposes is that even if *Strong Internalism* turned out to be true, this would only imply that every competent user of artifactual kind terms were in the same predicament as the epistemically most privileged user in the *Social Externalism* scenario. That is, while being extremely informed as to the relevant set of human intentions and purposes, they may still, *qua* competent users, have an ever so limited insight into many of the properties that may or may not make up actual instances of the kind in question. More importantly, they may, in the epistemic case, be ever so uninformed as for properties of epistemological *significance*—or so I will now argue.

First, consider the following definition:

Conceptual Accuracy.

A concept is *accurate* to the extent that it provides a *correct* and *complete* description of its referent.

The idea here is two-fold: (*a*) there are aspects of concepts that do not serve to determine reference, and (*b*) these aspects may be represented in terms of descriptions. Let us look closer at (*a*) first. Ever since Kripke (1980), it has been customary to distinguish between factors that serve to *fix* and factors that serve to *determine* reference. For example, while whatever conceptual component responsible for my tendency to think of horses as having four legs may serve to *fix* the reference of HORSE, i.e., to be a helpful tool in picking out actual horses in my environment, it does not *determine* reference, for the simple reason that some horses are amputees. As such, factors fixing reference, clearly, play an important cognitive role

in our mental life, by significantly facilitating our interaction with the extra-mental world. *Contra* the descriptivist, however, these factors should not be confused with the factors determining reference. We will return to this point in a second.

Let us now turn to (b). On closer inspection, we see that, if there were no conceptual aspects that could be represented in terms of descriptions, it is hard to see how concepts at all could be the objects of any kind of analysis in the first place. On any view of concepts—be it concepts as abilities, Forms, senses or mental representations—concepts serve to categorize the world and the systems of categorization that arise from concept use may be represented in terms of descriptions. Hence, a concept that serves to put all and only blue objects that weigh more than two pounds in one category may be characterized in terms of the description “is blue and weighs more than two pounds,” quite independently of one’s favored ontology of concepts. This is not, of course, to say that whatever mental event that is causally responsible for the categorization takes the form of a description—that would have to be established through empirical research—which is exactly why I am not saying that concepts *are* descriptions but merely that they may be *represented* as incorporating descriptions.

One particular reluctance to talking about concepts and descriptions in the same sentence stems from an aversion to the traditionally influential idea that reference determination works by way of the referent satisfying a description inherent in the corresponding concept. However, this is certainly not the idea being defended here. If anything, the present notion of conceptual accuracy serves to state in clearer terms the very externalism that served to refute this descriptivist picture of reference fixing: In so far as any form of content externalism holds, having an accurate concept—i.e. a concept whose descriptive aspect provides a correct and complete description of its referent—is *not* a prerequisite for successful reference. However, since we are, for the moment, assuming that *Strong Internalism* provides the correct semantics for artifactual kind terms, we will focus on the fact that not even *Strong Internalism* implies that having an accurate concept of an artifactual kind is a prerequisite for successfully referring to it. The reason may be brought out as follows:

Take any object x of epistemological investigation. If x is an artifactual kind, there is a set ϕ of human intentions and purposes that determines the identify conditions for x . What will ϕ contain? Given that x is an artifactual *epistemic* kind, it will most likely contain intentions and purposes pertaining to the attainment of certain epistemic goals, say, true belief in significant

matters. Is it possible to say something more specific? Well, if *Strong Internalism* holds, we may note that, unless we want to radically restrict the extent to which people may successfully refer to x , we have to restrict the richness of information contained in ϕ since, in the general case,

the more information is contained in ϕ , the rarer is successful reference to x .

Given that reference to x is widespread, however—which is, hopefully, the case for most epistemic kinds—we may, at the very least, say that

being acquainted with the information contained in ϕ can *not* involve a complete apprehension of all properties that make up actual instances of x .

Hence, even given *Strong Internalism*, the semantically most informed person—i.e., the person that has the most complete grasp of what is contained in ϕ —might still be in the dark as for many of the properties that may or may not make up instances of x . In other words, she may still have an inaccurate concept of the kind in question.

I take it that few would deny this claim, if understood in the weak sense of there always being further facts that could be found out that do not flow from what we know just by virtue of being able to successfully refer. For example, just by virtue of successfully referring to pens and keys, I may (at least on the *Strong Internalist's* story) know a whole host of things about pens and keys and, in particular, things that I may easily infer from being acquainted with the relevant sets of intentions and purposes. At the same time, there may very well be a lot of things that I do *not* know about pens and keys, such as its exact mechanical make-up, its molecular constitution, etc. As pointed out by Paul Griffiths (1997), this is to be expected considering that the factors determining the identity conditions are substantially less rich in the artifactual kind case than in the natural kind case:

The traditional natural kinds are among the richest. The kindhood of a physical element determines almost all its salient properties. [...] In contrast, knowing what sort of thing an artifact is, knowing that it is a bracelet for example, may fix very few of its features. There are just too many ways to skin a cat, or in this case too many ways to decoratively encircle the wrist. (Griffiths, 1997, 190)

Of course, it does not follow that being familiar with the multitude of ways that pens or keys may be crafted or wrists may be decoratively encircled is necessarily very significant to me, given that my goals, as far as pens, keys, and bracelets go, are restricted to successful, everyday interactions. What I want to claim, however, is that the same does *not* hold in the epistemological case. In particular, I want to claim that what is not contained in ϕ —i.e., what is still to be found out when we have enough knowledge for successful reference to occur—is of *epistemological significance*.

My argument for this claim runs as follows. First, consider the following: One important epistemological desideratum is to *guide* epistemic inquiry. One element in this desideratum is to see to it that our epistemic vocabulary is as *apt* as possible, where a vocabulary is apt in so far as it invokes apt concepts, and

Conceptual Aptness.

a concept is *apt* to the extent that it serves its intended purpose well.

To say that concepts serve purposes is not meant to imply anything controversial. At a very basic level, the purpose of concepts is simply to enable us to think certain thoughts, have certain beliefs, etc., and, thereby, interact with the world in more or less successful ways. Differently put, concepts provide a *framework* for thinking and believing (and so on), in the sense of a way of categorizing the world. As such, the extent to which one framework is better than another depends on what we wish to attain with it.

This point may be illustrated in the epistemic domain by noting that we, as epistemic inquirers, are engaged in a certain project of epistemic evaluation and doxastic revision, (roughly) aimed at attaining and maintaining true belief in significant matters. An integral part of succeeding in this latter task is having an apt epistemic vocabulary, where an apt epistemic vocabulary is a vocabulary that can be used to categorize and, thereby, evaluate fellow inquirers and the world in a way that serves the purpose of attaining and maintaining true beliefs in significant matters. Clearly, some concepts will serve this purpose better than others. In particular, the following seems a reasonable claim:

If epistemic vocabulary V_1 is more *refined* than vocabulary V_2 —i.e., if V_1 incorporates *accurate* concepts to a larger extent than V_2 does—then V_1 is more *apt* than V_2 , *ceteris paribus*.

Furthermore, it is reasonable to assume that getting acquainted with the properties that may or may not make up actual instances of x —and, in particular, those properties that may not be readily inferred from being acquainted with what is contained in ϕ —would enable us to refine our concept of x , in the sense of pruning it so as to provide a more correct and complete description of its referent. Finally, it seems fairly uncontroversial that the proper method for getting acquainted with those properties will have to be empirical.

Take justification, by way of illustration. Construed as an artifactual kind, the relevant set of intentions and purposes would most likely pertain to the flagging of appropriate sources of information, where the appropriateness is understood in terms of truth-conductivity. More than that, on *Strong Internalism*, every competent speaker would be perfectly familiar with the details of this set. But does this imply that they, thereby, know everything there is to know about justification? Does it, in particular, follow that there is nothing else to find out that is of *epistemological significance*? That seems unreasonable. In particular, it would be of epistemological significance to find out, among other things, (a) what external phenomena actually satisfy the relevant requirements of truth-conductivity, (b) about the multiplicity of properties that may or may not make up these phenomena, and, perhaps more importantly, (c) how they fit in to the causal fabric of the world and, hence, may not only be better understood but also be manipulated to the benefit of the epistemic inquirer—all of which seems to be things that cannot in any straightforward way be inferred from the relevant set of intentions and purposes, nor be discovered without recourse to an empirical investigation.

So, in the general case, and in so far as conceptual *refinement* may rightfully play a substantial role in epistemology, the claim that epistemological investigation is a substantially empirical investigation is largely independent not only of whether the objects of epistemological investigations are natural or artifactual kinds, but also of any content externalist or internalist considerations with respect to the latter. In other words, given that the above line of reasoning does, indeed, apply to most (if not all) objects of epistemological investigation—and I see no reason why it would not—we may conclude that the claim that epistemological analysis is substantially empirical does not hinge on (A). Indeed, the preceding discussion provides us with reason to take the following argument to be sound:

- (A*) For every object x of epistemological investigation, x is either a natural or an artifactual kind.
- (B*) If (A*), epistemological investigation is substantially empirical.
- (C) Hence, epistemological investigation is substantially empirical (A*, B*, *MP*).

This concludes my solution to *Problem 1*. Next, we will be looking into why (D) does *not* hold—i.e., why it does not follow from (C) that an understanding of our epistemic concepts is largely irrelevant to epistemological investigation. In the process of doing so, we will not only provide further evidence to the effect that epistemology is substantially empirical, but also introduce a more radical means to attaining an apt vocabulary: conceptual reconstruction.

5. *An alleged rationale for (D): Factual analysis and the AC principle*

Returning to *The Argument*, let us now turn to premise (D) and the conclusion (E), stating that a thorough understanding of our epistemic concepts, over and above the phenomena that they pick out, is irrelevant to epistemological investigation. Kornblith's commitment to this conclusion comes out most clearly in his critique of the idea that the (sole) job of epistemology is to analyze epistemic concepts by way of intuitions about hypothetical cases—a view that he characterizes as follows:

Appeals to intuition are designed to allow us to illuminate the contours of our concepts. By examining our intuitions about imaginary or hypothetical cases, we should be able to come to an understanding of our concepts of, for example, knowledge and justification. The goal of epistemology on this view, or, at a minimum, an essential first step in developing an epistemological theory, is an understanding of our concepts. (Kornblith, 2006, 11–12)

Kornblith continues:

My own view is that our concepts of knowledge and justification are of no epistemological interest. The proper objects of epistemological theorizing are knowledge and justification themselves, rather than our concepts of them. (Kornblith, 2006, 12)

In light of the reasonable claim that *some* initial examination of our epistemic concepts might be necessary in order to fix the subject matter,

Kornblith makes it clear that his main disagreement with the tradition of epistemology as conceptual analysis concerns the scope of such a semantic investigation. More specifically, he claims that the semantic investigation called for is “utterly trivial” and, thereby, in no way related to the two thousand year old project that, in a tradition stemming from Plato’s *theaetetus* and culminating in the Gettier-inspired literature, typically falls under the heading of the analysis of knowledge and justification.

When trying to put this claim in more precise terms, it serves us well to make a distinction between two stages of epistemological investigation. The first one corresponds to the *identification* of an epistemological object *F*, i.e., of fixing the subject matter (if only tentatively) through picking out a selection of what we take to be paradigmatic instances of the concept ‘*F*.’ In doing this, our concepts ‘*F*’ (and the categorization intuitions they give rise to), clearly, play a vital role, if only in the sense that, in order to identify an *F*, one needs to have some kind of grasp of what it is to be (an) *F*. Naturally, depending on the extent to which ‘*F*’ is ambiguous or imprecise, this process of identification may be more or less demanding. Regardless, the purpose of identification is to pave the way for the more substantial and straightforwardly empirical *aggregation* of the characteristics that may or may not be found in *F*s, in order to reach a satisfactory answer to the question that incited inquiry in the first place: “What is (an) *F*?”

Against the background of this distinction, I would like to characterize Kornblith’s notion of analysis as follows:

Factual Analysis (FA).

Identification: For any concept ‘*F*,’ Identify a set *Q*, containing a selection of what we take to be paradigmatic instances of ‘*F*.’

Aggregation: Against the background of an *a posteriori* investigation into the (extra-mental) elements found in *Q*, aggregate a set of characteristics that specify what actually constitutes (an) *F*.

Is FA an empirical of analysis? In so far as aggregation goes, the answer would have to be ‘yes.’ As stated, however, it remains an open question whether the same goes for identification. More specifically, the question

9. This distinction between *identification* and *aggregation* is borrowed from Sen’s (1981) treatment of the issue of poverty.

remains open until it is specified whether the fixing of a set of paradigmatic examples presupposes an *a priori* access to conceptual content as opposed to an empirical investigations into what people tend to mean by terms. However, the question does not need to be settled for present purposes since it seems reasonable to assume that either of the two claims would have to be true: Either (a) identification calls for a substantial investigation, in which case an empirical method will be desirable due to a superior methodological rigor, or (b) identification does *not* call for a substantial investigation, in which case an *a priori* method would not compromise the claim that epistemological investigation—if understood along the lines of FA—is still a *substantially* empirical investigation.

We have to keep in mind, however, that Kornblith does not only commit himself to the idea that FA as a *viable* method of epistemological analysis, but to the stronger claim that FA provides a *complete* method, i.e., that there are no other aspects to epistemological investigation over and above identification and aggregation, as spelled out above. This brings us to (D). If FA is all there is to epistemological analysis, there is no need for a thorough understanding of our epistemic concepts—as in an understanding that goes beyond whatever semantic investigation is needed for identification—since the main component of epistemological analysis will consist in a purely empirical investigation into the epistemic phenomena picked out. In other words, if FA provides a complete methodology, we should accept (D).

But why should we take FA to provide a complete methodology? Suppose that it proves possible to successfully conduct a series of FAs, providing an account of what constitutes (an) *F* for *every* epistemic concept '*F*.' This would, undoubtedly, be an impressive accomplishment. But would it mark the end of epistemological investigation? Considering the picture of epistemology as the pursuit of an apt epistemic vocabulary, the question may be reformulated as follows: Would such a set of analyses necessarily yield a fully apt epistemic vocabulary? Considering what was said above in relation to conceptual refinement, it might be tempting to answer the question in the positive, under the assumption that a more apt set of concept simply is a more *accurate* set of concepts. In other words, the answer would be 'yes' if the following principle can be shown to hold:

The AC principle.

The pursuit of a more *apt* set of concepts reduces to that of providing a more *accurate* set of concepts.

In other words, we have established the following chain of dependency: If *the AC Principle* holds, then it is reasonable to assume that FA yields not only accurate but apt concepts and, hence, provides a complete epistemological methodology. Furthermore, if FA provides a complete epistemology methodology, then it is reasonable to assume that (D) holds. However, if *the AC Principle* does *not* hold, (D) remains unwarranted since FA, thereby, might yield accurate but not necessarily apt concepts. More specifically, I will argue that

Problem 2.

unless *the AC Principle* holds, FA fails to acknowledge epistemology's role in the particular kind of conceptual improvement involved in conceptual *reconstruction*.

In the following two sections, I will (a) spell out this methodological component of conceptual reconstruction, (b) provide two examples of cases in which the need for it indicates that increased accuracy does not imply increased aptness, even if we assume content externalism for the corresponding concepts, and (c) conclude that *the AC Principle* is an unviable epistemological assumption and (D) is without warrant.

6. *Conceptual aptness in epistemology*

A strong motivation for content externalism about natural kind terms is that it provides us with a straightforward and intuitive explanation of disagreement in the natural sciences.¹⁰ That is, given content externalism, successful reference is completely compatible with inaccurate concepts and theories on part of the person referring. Hence, even if Dalton, Rutherford, and modern physicists were and are working with substantially different theories and concepts of the atom, they were and are essentially talking about the same phenomenon, i.e., the atom. Indeed, only if we say this can we make the further claim that they *disagree* about the latter's constitution and that contemporary theories of the atom constitute *improvements* of the earlier theories of Dalton and Rutherford, in line with the overall progress of science.

10. See, e.g., Putnam (1975).

Does this rather neat picture of scientific disagreement and improvement carry over to epistemology? One of the major reasons for doubting that it does is that (a) we have reason to believe that the majority of epistemological objects are artifactual rather than natural kinds, that (b) it remains to be shown that content externalism provides the correct semantics for artifactual kind terms, and that (c) it is not obvious that explaining epistemological disagreement requires assuming *referential* continuity, rather than that our epistemological *project* is continuous with that of our epistemological predecessors (e.g., in the sense that we are all concerned with the search for an apt epistemic vocabulary). However, being able to explain disagreement is, clearly, only one motivation for content externalism and there might very well be independent reasons for wanting to defend such externalism in epistemology. For this reason, I will now consider a dialectically more robust strategy by assuming content externalism but showing that *the AC principle* still does not hold. The strategy will first be demonstrated on an abstract level and then, in the next section, illustrated by way of two examples.

So, first consider how referents get assigned to concepts on an externalist story. For dialectical purposes, let us assume the strongest form of externalism, i.e., *Causal Externalism*. On *Causal Externalism*, a term or concept gets assigned a referent by way of an initial act of baptism. In other words, the referent of a concept is determined by the (mere) fact that the baptizer stands in a certain causal relation to it. As we have already seen, such a semantic story is perfectly compatible with successful reference in spite of considerable ignorance on part of the speaker, which leaves room for extensive conceptual refinement. Hence, it was suggested above that there is a connection between refinement and aptness. This, furthermore, provides at least part of a rationale for the conceptual refinement of purely descriptive concepts in science, understood as concepts the mere (or at least most central) purpose of which is to categorize, without thereby providing a normative evaluation of whether something is good or bad in relation to a specified set of goals. Hence, H₂O may constitute a refinement of water, as far as chemistry goes, and mean molecular kinetic energy a refinement of temperature, as far as kinetic theory goes.

However, to fully understand what constitutes an apt *epistemic* vocabulary, we need to keep in mind that epistemic concepts typically are *normative* concepts. Epistemic concepts, *qua* normative concepts, are tools by which we evaluate the conduct of fellow epistemic inquirers—that is, categorize their conduct as being an instance of knowledge, justification,

rationality, etc.—but where we, merely by virtue of such a categorization, also make an explicitly normative judgment about the extent to which what is being done is good from an epistemic point of view. This suggests that the purposes of epistemic concepts may be understood in relation to the *principles* in which they figure and the *goals* that these principles are designed to meet. Furthermore, it prompts the following question: Is there any guarantee that the *referent* that was originally fixed to an epistemic concept or term in an initial act of baptism, in fact, provides the best route to our epistemic goals?

The answer is no. Granted, it might be possible to construct an argument to the effect that whatever we are referring to with our epistemic concepts does not provide a completely *useless* route to our epistemic goals, given that having true belief is an integral part of attaining many of our practical goals (and that the latter is something that we tend to do). However, it is hard to see why the referent initially fixed necessarily provides the *best* path to our epistemic goals. For this reason, epistemology has to consider the possibility that alternative referents may present a better route to our epistemic goals than our present referents. Hence, it makes sense to not only investigate the question of to what extent our concepts provide *accurate* pictures of their referents—referents that need not provide the best route to our epistemic goal—but also whether there are any *alternative* referents that might present a better route. Such an inquiry, however, has to take into account not only the causal structure of the world (as revealed through straightforward empirical inquiry) but also the purpose of the original concepts and, in particular, the principles in which they typically figure and the goals these principles are meant to attain, since nothing short of such an investigation will enable us to specify what would constitute a better route to our goals and, consequently, a more apt concept.

7. *When accuracy does not increase aptness: Two scenarios*

To illustrate this point, I will now present two hypothetical scenarios in which increased accuracy does not increase aptness, since the referents in question—fixed in accordance with *Causal Externalism*—do not present optimal routes to our epistemic goals. I will argue that the proper epistemological strategy in those cases is not refinement but a more substantive conceptual improvement in light of the larger context of principles and

goals in which the concept figures. I will refer to such conceptual improvement as conceptual *reconstruction*, to a first approximation understood as an ameliorative activity located further out on a continuum of increasingly radical conceptual revision.

A helpful metaphor here is architectural reconstruction. When reconstructing, say, a house, you start out with certain pre-existing material, i.e., the house that is to be reconstructed. Let us call this house H_1 . The house that results from the reconstruction—let us call it H_2 —might look nothing like H_1 . Nevertheless, H_2 will (if everything goes as planned) serve a certain set of purposes better than H_1 did. Perhaps H_2 is more spacious, has better insulation, has a more attractive design, etc., than H_1 . Indeed, it is reasonable to believe that the intention to realize those properties provided the very reason for reconstruction. Similarly, conceptual reconstruction starts out with a pre-existing concept. The reconstructed concept might, in the end, look nothing like the original concept. Still, the very point of reconstruction is that the reconstructed concept serves a set of purposes better than the original concept. In the case of epistemic concepts, these purposes will be understood in relation to our epistemic goals—goals that will be spelled out in more detail below but that, to a first approximation, may be understood in terms of true belief in significant matters.

However, this raises some questions about the identity conditions for concepts. Is the concept that results from reconstruction “the same concept” as the concept that we started out with? I find this question about as puzzling (and interesting) as the question whether my kitchen remains “the same” over the course of a renovation. Clearly, many of its properties will change—some vanish, some arise—and the reconstructed kitchen will not be identical to the old one. (If it were, the renovation would have failed.) At the same time, it is still my kitchen and it will still serve the same purposes—indeed, it will, hopefully, serve some of the same purposes *better* than my old kitchen. Analogously, I will say that a reconstructed concept C_2 is “the same” concept as an original concept C_1 to the extent that they both figure in relation to the same set of purposes. At the same time, however, C_2 will, clearly, also be different from C_1 in the sense that it has different properties and, due to this, serves the purposes in question to a greater degree.

Against the background of a distinction between conceptual refinement and reconstruction, say that we perform a factual analysis of JUSTIFICATION and let us, for simplicity’s sake, refer to this concept as *our*

concept of justification.¹¹ Assume, furthermore, that we, in the process of identification, find that the properties by which we typically individuate degree of justification pertain to the fulfillment of epistemic duties. In fact, this would make complete sense, given what we know about the etymology of the term “justification.” As noted by William Alston, the term “has been imported into epistemology from talk about voluntary action,” which “explains the strong tendency to think of the justification of belief in deontological terms, in terms of being permitted to believe that *p* (not being to blame for doing so, being ‘in the clear’ in so believing)” (Alston, 1993, pp. 532 and 533, respectively).¹² What does this tell us about the way that JUSTIFICATION was originally endowed with a referent? For one thing, it lends some support to the dual claim that (a) JUSTIFICATION was originally introduced as applying to the formation of belief, and that (b) it has traditionally been presupposed that we can form or refrain from forming beliefs by willed action—on pain of denying that ought implies can. Finally, assume that we find, in the process of aggregation and empirical investigation of the phenomenon actually picked out by our concepts, that we have no voluntary control over the formation of beliefs.¹³

If this turned out to be the case, we seem to have uncovered reason to believe that our concept of justification is off the mark, in that it pertains to something that we, as a matter of fact, cannot have, namely epistemic duties. What would be a proper epistemological response? Two responses are available. On the first response, we reject voluntarism as a mistaken view about the way our mind works, but retain the idea that justification applies to belief-formation. This would correspond to a simple refinement and the most promising candidate for fleshing it out would probably be some form of process reliabilism. However, as it stands, this response

11. It does not matter so much for present purposes whether there is such a thing as our concept of justification or rather a rich multiplicity, since all that the following line of reasoning requires is that we are talking in terms of a specific concept—be it a widely shared one or not.

12. See also Plantinga (1990).

13. The kind of voluntary control at issue here is what Alston (2005, 62) refers to as *basic voluntary control*. Undoubtedly, this is not the only kind of voluntary control—in fact, Alston distinguishes between three types of (decreasingly extensive) voluntary control as well as different grades of indirect voluntary *influence* (pp. 62–80). However, Alston also provides convincing arguments to the effect that, even given increased taxonomical complexity, there does not seem to be any voluntary control or influence such that it both (a) applies to the psychology of common epistemic inquirers, and (b) is sufficiently extensive to warrant talking in terms of genuine duties and obligations.

suffers from a significant problem: It takes for granted that the referent inherited from our deontological predecessors does, in fact, provide the optimal route to our epistemic goals. While this certainly cannot be ruled out, nor can it be assumed, for reasons brought out in the previous section.

This leads us to the second response, on which we engage in an inquiry best described as a continuation of aggregation, with the crucial qualification that it is preceded and guided by an investigation into the intended purpose of the original concept, in an empirical search for properties that may serve that purpose better, given a relevant set of epistemic principles and goals. For example, if such an investigation were to demonstrate that the purpose of JUSTIFICATION is to flag certain voluntary acts (previously identified as acts of belief-formation) as appropriate sources of information, and the appropriateness in question is typically understood in relation to a goal of attaining and maintaining true belief in significant matters, one possible route for empirical inquiry would be to identify a kind of voluntary act that tends to yield and support true belief, thereby providing material for a reconstructed concept. The resulting view would retain voluntarism but reject the idea that justification should apply to belief-formation. This would correspond not to a refinement but a reconstruction, where justification gets “re-baptized,” so to speak, and JUSTIFICATION, thereby, gets assigned a new referent.

This brings us to the second scenario and the traditionally most influential candidate for such a voluntary act: *introspection*. More specifically, say that we, in the process of identification, find that we tend to determine degree of justification by reference to an introspective evaluation of reasons on part of the allegedly justified (or unjustified) subject. Let us, furthermore, assume that this particular notion of JUSTIFICATION originally entered into the discourse of evaluating epistemic subjects some four hundred years ago by way of Descartes’ ideas about what one perceives clearly and distinctly on introspection. What does this tell us about the way in which JUSTIFICATION was originally endowed with a referent? At the very least, it lends some support to the dual claim that (a) JUSTIFICATION was originally introduced as applying to acts of introspection, and that (b) it has traditionally been presupposed that such acts provide a powerful and reliable access to the grounds for our beliefs. However, suppose that we also find that we, as a matter of empirical fact (say, facts uncovered by cognitive psychology), seldom have access to the epistemic qualities of the processes by which we form beliefs and, furthermore, that the stories (consciously

or unconsciously) reconstructed by us as to the epistemic etiology of our beliefs are often quite inaccurate.¹⁴

If that turned out to be the case, what would be a proper epistemological response? Again, two responses are available. On the first response, we would try to identify conditions under which we *do* have reliable access to our reasons and, then, refine our concept accordingly. However, the very same research hinted at in the previous paragraph gives us reason to think that such conditions are quite hard to come by. Hence, the second response: Conduct an investigation into the purpose of our (supposedly inapt) concept, let the result of such an investigation guide further empirical aggregation of candidate properties that may figure in a reconstructed concept that fills the same (or close to the same) purpose, without being committed to the idea that we have a reliable introspective access to the epistemic qualities of our belief-forming processes. For example, if an investigation into the purpose of our concept of justification were to reveal that its purpose is to flag certain voluntary acts (previously identified as acts of introspection) as appropriate sources of information, and the appropriateness is (again) typically understood in relation to a goal of attaining and maintaining true belief in significant matters, one way for empirical inquiry to proceed would be to identify an alternative kind of voluntary act (one candidate being certain acts of reasoning) that tends to yield and support true belief. An empirical aggregation preceded and guided by an investigation into the purpose of our concept would, thereby, provide material for a reconstructed concept.

These hypothetical examples serve to illustrate two points with a direct bearing on (D) and the issue of accuracy and aptness: First, there are possible cases in which merely attending to the referents of our current concepts would *not* enable us to complete the task of identifying more apt concepts, since those referents provide sub-optimal routes to our epistemic goals. For this reason, an epistemological investigation need attend not only to the referents of our concepts but also to other properties that are not in any obvious way implicated by our current concepts but that may, nevertheless, figure in a more apt vocabulary. Second, this empirical investigation needs to be preceded and guided by an investigation into the purposes of the original concepts, providing the empirical inquiry at issue with a direction in the form of an understanding of what properties to look for. For this reason, the substantially empirical method in no way

14. See Wilson (2002) for some evidence to this effect.

eliminates the need for an understanding of our epistemic concepts and, in particular, the particular purposes for which we employ them. To the contrary, such an understanding plays a vital role in the search for a more apt epistemic vocabulary.

Unless it can, somehow, be shown that scenarios like the two just considered are impossible (which is different from arguing that they are not actual), we have to leave room for the possibility of conceptual reconstruction when providing a methodological framework for epistemology. More specifically: It is beyond doubt that some instances of conceptual improvement in epistemology might indeed flow from a straightforward conceptual refinement. However, given that epistemology is concerned with explicitly *normative* concepts, some cases of improvement must take into account not only (*a*) facts about the referent but also (*b*) facts about the intended *purpose* of the concept in question, so as to guide empirical aggregation in the search for (*c*) properties that might not be implied by the original concept, nor present themselves in any straightforward way through an unconditional empirical investigation into the referent, but that might nevertheless furnish a reconstructed concept with an increased aptness, given the principles and goals that the original concept was supposed to (but failed to) meet. Hence, even if substantially empirical, a proper epistemological methodology needs to leave room for attending to our concepts, and in particular to the purpose for which we employ the epistemic concepts that we do. This is why epistemic concepts—*contra* the decrees of FA—do not drop out of the epistemological picture as soon as we move beyond the initial stage of delimiting a set of paradigmatic examples.

8. *Conclusions*

By way of recapitulation, we have established two conclusions. The first conclusion is that (C) may be established independently of (A). That is, under the plausible assumption that conceptual *refinement* plays an important role in epistemological theorizing, the claim that epistemology is a substantially empirical investigation is *not* contingent upon the admittedly controversial idea that all objects of epistemological investigation are natural kinds. In fact, (C) can be shown to be plausible even under the dual assumption that (*a*) all objects of epistemological investigation are *artifactual* kinds, and that (*b*) what I have referred to as *Strong Internalism*

provides the correct semantics for artifactual kind terms. The second conclusion is that (E) does not follow from (C), since (D) is false. That is, concepts are not only relevant to identification, i.e., the fixing of a (non-exhaustive) set of (what we take to be) paradigmatic examples of the phenomenon under investigation. An insight into the purposes of our epistemic concepts is, in some cases of conceptual reconstruction, also a prerequisite for knowing how to direct the process of aggregation in the improvement of our conceptual apparatus and, hence, answering a question that ought to lie at the heart of any epistemology interested in not only describing but also improving on epistemic inquiry, namely “Given our epistemic goals, what would be a set of epistemic concepts that *served us better?*”

REFERENCES

- Alston, W. P., (2005), *Beyond “Justification”: Dimensions of Epistemic Evaluation*, Cornell University Press.
- Burge, T., (1986), “Individualism and Psychology,” *Philosophical Review* 95, 3–45.
- Goldman, A., and Pust, J., (1998), “Philosophical Theory and Intuitional Evidence,” in M. DePaul and W. Ramsey, eds., *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, Inc., 179–197.
- Griffiths, P. E., (1997), *What Emotions Really Are: The Problem of Psychological Categories*, The University of Chicago Press.
- Kornblith, H., (1993), *Inductive Inference and Its Natural Ground*, The MIT Press.
- (2002), *Knowledge and Its Place in Nature*, Oxford University Press.
- (2006), “Appeals to Intuition and the Ambitions of Epistemology,” in S. Hetherington, ed., *Epistemology Futures*, Oxford University Press, 10–25.
- (2007), “Naturalism and Intuitions,” *Grazer Philosophische Studien* 74, 27–49.
- (forthcoming), “How to Refer to Artifacts,” in E. Margolis and S. Laurence, eds., *Creations of the Mind: Essays on Artifacts and their Representation*, Oxford University Press, forthcoming.
- Plantinga, A., (1990), “Justification in the 20th Century,” *Philosophy and Phenomenological Research* 50 (supplement), 45–71.
- Putnam, H., (1975), “The Meaning of ‘Meaning’,” reprinted in his *Mind, Language and Reality: Philosophical Papers* vol. 2, Cambridge University Press, 215–271.

- Sen, A., (1981), *Poverty and Famines: An Essay on Entitlement and Deprivation*, Oxford University Press.
- Thomasson, A. L., (2003), "Realism and Human Kinds," *Philosophy and Phenomenological Research* LXVII (3), 580–609.
- Wilson, T., (2002), *Strangers To Ourselves: Discovering the Adaptive Unconscious*, Harvard University Press.