15

How Much Does Knowledge Matter? Reality, Representation, and the Aims of Epistemology

TRISTRAM MCPHERSON
AND DAVID PLUNKETT

1 Introduction

Much of philosophical inquiry about epistemology aims ultimately to understand epistemic properties, relations, and facts themselves, as opposed to our epistemic thought and talk about them. For example, consider the following familiar epistemological questions: What does knowledge consist in? When (and why) is someone epistemically justified in believing a proposition? When is it epistemically rational to believe a proposition based on testimony? These appear to be *object-level* questions within epistemology: questions about what we might call "epistemic reality" itself.[1] These questions contrast with *representational-level* questions about our epistemic thought and talk (e.g., about the meaning of the word 'knowledge', or how to analyze the concept RATIONAL).[2]

Suppose one aims to engage in object-level inquiry in epistemology. This endeavor leaves open the following question: *which* epistemic entities should an object-level epistemic inquirer investigate? (For brevity, here and in what follows, we use the phrase 'epistemic entities' broadly, to pick out those parts of "epistemic reality"—such as epistemic properties, relations, kinds, and facts—that matter for our discussion.) One answer is that we should investigate those entities picked out by our pre-theoretical words and concepts: the kinds of words and concepts that regular people ("the folk") use in everyday life. This kind of view seems to secure an important good for epistemic inquiry: namely, it seems to ensure that when epistemologists put forward their philosophical theories (e.g., about "knowledge" or "justification"), they are discussing things that motivate turning to epistemological inquiry to begin with, rather than "changing the subject" to something else.[3]

In this chapter, we advocate for a competitor to the view just sketched. We argue that object-level epistemologists should investigate those epistemic entities that are apt targets for

Kornblith and His Critics, First Edition. Edited by Joshua DiPaolo and Luis R.G. Oliveira.
© 2026 John Wiley & Sons, Inc. Published 2026 by John Wiley & Sons, Inc.

serious inquiry. What makes them apt, we claim, is primarily a matter of the explanatory and normative importance of those entities, rather than something about their relationship to our words and concepts (whether "pre-theoretical" ones or not). After introducing and motivating this view in general terms, we illustrate its significance, and how it contrasts with the view sketched in the preceding paragraph, by offering a case study: Hilary Kornblith's work in *Knowledge and Its Place in Nature*.[4] Among contemporary epistemologists, Kornblith has been one of the most influential proponents of a thoroughly object-level focus for epistemic theorizing, and one of the strongest critics of focusing on issues about ordinary epistemic thought and talk. Against this background, we argue that Kornblith can (and given his own commitments, *should*) be more thoroughgoing in detaching his epistemic inquiry from ordinary epistemic thought and talk. Kornblith's target in his book is *knowledge*, which he understands to be the relation picked out by our pre-theoretical "knowledge" thought and talk. We argue that even if Kornblith has offered a completely successful account of this relation, he could have failed to locate and study the most apt targets for epistemic inquiry.

We proceed as follows. First, we set out the two goals that structure our proposed approach to "object-level" epistemic inquiry. These goals reflect our *explanatory* and *normative* interests in epistemology (Section 2). We then briefly introduce some of Kornblith's key commitments in *Knowledge and Its Place in Nature* (Section 3). With this context in place, we critique Kornblith's approach, relative to the explanatory and normative goals we have introduced, respectively (Sections 4 and 5). Finally, we explore what role inquiry into epistemic representations should play within object-level epistemic inquiry driven by our two goals (Section 6). Here, we draw on recent work in "conceptual ethics" and "conceptual engineering." Put roughly, we argue that the approach to object-level inquiry (whether in epistemology or elsewhere) that we advertise here can motivate, or even require, engaging in normative and evaluative assessment of our concepts and words. In turn, this might well lead to work on improving or replacing those concepts and words.

Through much of the chapter, we grant Kornblith many controversial commitments, and thus develop our argument as something close to an internal critique. In the conclusion, we step back from this approach, to locate our argument in this chapter within our own broader framework for thinking about inquiry in epistemology. We emphasize how the "object-level" inquiry that we advertise here is only one of multiple distinct forms of epistemic inquiry that we think are important, as well as briefly discuss some of the ways in which these forms of epistemic inquiry can inform each other.

While we take Kornblith's pathbreaking work as our foil, the methodological issues for "object-level" theorizing that we illustrate in this chapter are of much wider significance, across both epistemology and many other subfields of philosophy. Our hope is that this paper can help to bring these issues more clearly into view, and to illustrate their significance.

2 The Goals of Object-Level Epistemic Inquiry

The core idea of the view we advocate for in this chapter is this: as epistemologists, we should inquire into those epistemic entities, *whatever they are*, that best enable us to achieve the central goals of epistemic inquiry. The question of what the central goals of

epistemic inquiry are (or, perhaps more importantly, should be) is controversial. For example, one might think that epistemic inquiry could be instrumentally valuable for promoting a more just or emancipated society, or for promoting an ongoing conversation, of a certain sort, in our social/historical tradition.[5] For the purposes of this chapter, we focus on two goals that we think animate many parts of contemporary philosophical work on epistemology. The first is to learn about epistemic entities that are *explanatorily* important. The second is to learn about ones that are *normatively* important. In this section, we explain and motivate these goals. In the sections that follow, we illustrate how pursuing these goals might lead us to investigate entities other than those picked out by our pretheoretical conceptual and linguistic schemes.[6]

Start with the goal that involves explanatory importance. The core idea here is that we should inquire into those epistemic entities that play robust roles in explaining descriptive facts. One way to motivate this goal is that a range of epistemic entities seem to (to some degree) help explain descriptive facts, such as facts about what people actually do, say, or think. For example, it is common to appeal to facts about what people know, or have evidence for believing, in explaining their actions. Where epistemic entities play such roles, it is natural to think that epistemic inquiry should illuminate facts about such entities, and their explanatory roles. Another way to motivate this goal is to look to domains that we take to be models of successful inquiry. A familiar and plausible thought is that some of our most successful areas of inquiry—including key parts of the natural and social sciences—aim to identify and illuminate entities that do genuine explanatory work. With this thought in mind, one might seek to model epistemic inquiry on these successful forms of inquiry.

This goal has two concrete implications for epistemic inquiry.

First, we can ask a "domain-general" question about what sorts of entities are generally fit to play explanatory roles. Here, we can distinguish two ways of understanding the idea of "explanation." On the one hand, we can understand it in broadly "epistemic" terms, where (roughly) the success-conditions of explanation involve helping agents to better understand an issue. (Think here of "explaining" to someone where to find a parking lot to park their car.) On the other hand, we can understand "explanation" in broadly "metaphysical" terms, where (roughly) the success-conditions of explanation involve accurate representation of the actual facts about why things happen, or about what constitutes what. For our purposes here, we focus on this second understanding of explanation. Here, it's plausible that, whether in epistemology or elsewhere, entities that constitute "joints of nature" or "natural kinds" are apt to do such explanatory work, while the mere fact that we happen to think and talk about an entity does not ensure that it can do such work.

Second, we can focus specifically on the sorts of explanatory roles likely played by epistemic entities in particular. Notably, epistemic entities appear to play a series of distinctive roles in our lives: such as in inquiry, belief management, assertion, and action. So, other things being equal, identifying entities that are apt to play these distinctive roles in particular will be an especially important aspect of achieving this first goal.

Now turn to the second goal, about normative importance. The core idea here is twofold. First, we should inquire into those epistemic entities that play robust normative or evaluative roles. Second, we should seek to understand the ways in which these entities are

normative or evaluative. (In this chapter, we often use the term 'normative' broadly, to cover both the normative and the evaluative.) One way to motivate this goal is that epistemic entities seem to play several notable kinds of normative or evaluative roles. For example, consider the following familiar and plausible claims: knowledge *licenses* assertion; evidence constitutes *normative reasons* to believe; and epistemically irrational beliefs can be *blameworthy*. This goal can also be motivated on more general grounds. We take the (apparent) normativity of the epistemic to partly explain our interest in epistemic inquiry. Thus, just as one might hope that ethical inquiry can help us to live well, so we might hope that epistemic inquiry could help us to better regulate ourselves in relevant ways. For example, if we assume that it is good to be properly guided by knowledge, then inquiring into what knowledge is will potentially enable us to be more properly guided by knowledge, either individually or collectively (e.g., via our thought and talk about it, our practices around it, etc.).[7]

It is of course possible to dispute the significance of one or both of these goals in epistemic inquiry. For example: one might think that epistemic entities (or at least many important ones) only play a minor explanatory role in our best descriptive accounts of behavior etc., or one might try to debunk (rather than explain) their apparent normativity. However, we take these goals to animate much epistemic inquiry, both historically and in the contemporary field. In this chapter, we will simply assume that they are apt goals, and sketch some methodological lessons that follow from taking them seriously.

If one takes these goals seriously, one should be especially interested in the epistemic entities that are *most* explanatorily and/or normatively important. This is one of the core claims we advocate for in what follows. Doing so will allow us to make vivid how our understanding of the aims of epistemic inquiry can matter, and how our understanding of such inquiry contrasts with approaches that tether it more closely to facts about our pretheoretical concepts and words. However, it should be stressed that studying epistemic entities that are explanatorily or normatively important *to some degree*, *along certain dimensions* might well matter a lot to our overall goals here, even if there are many other ones that are *more* important. Consider an analogy: many different parts of a car engine help explain how it works, even if there are some that are more important than others.

3 Kornblith on Investigating Knowledge

As we read him, Kornblith is sensitive to both of the goals just introduced. Indeed, we take his work to be one of the leading models of systematic epistemic inquiry that combines concern for each of these goals with a robustly naturalistic methodology. To make clear how his work combines these goals—and to set the stage for our broader argument—we now introduce several key aspects of Kornblith's work on epistemology. For the sake of concreteness, and because of its centrality to Kornblith's work overall, we focus on the account of knowledge that Kornblith develops in *Knowledge and Its Place in Nature*. In this section, we sketch the following six aspects of Kornblith's views: his advocacy for an object-level focus for epistemic inquiry; his idea that knowledge is a natural kind; his proposed naturalistic methodology for learning about that kind; his substantive account

of the nature of that kind; his account of the normative significance of that kind; and the externalist metasemantics that underlies his project.

Kornblith begins *Knowledge and Its Place in Nature* with a set of bold methodological claims. These claims both set out his own positive agenda and provide the basis for criticizing others' approaches to epistemology. He writes that "the subject matter of epistemology is knowledge itself, not our concept of knowledge."[8] So, according to Kornblith, epistemic inquiry is, in the terms we introduced in the introduction, a kind of object-level inquiry. One might think that, even if one's ultimate goal is to understand some object-level phenomenon, there is a crucial role for representational-level inquiry. For example: it might help one better orient to, or understand the boundaries of, the basic subject matter one is trying to investigate.[9] But, for Kornblith, inquiry into our representations does not play even a significant subsidiary role of this kind. Because of this belief, he boldly claims that "epistemology should not concern itself with our concept of knowledge."[10]

Kornblith takes knowledge to be a *natural kind*. On his account, a natural kind is a cluster of properties that stand in homeostatic relations: roughly, relations of nomically stable covariation and causation.[11] This underlying stability explains the ability of natural kinds to underwrite strong inductive inferences and causal explanations. And this ability, Kornblith claims, in turn explains why scientific inquiry tends to be organized around natural kinds.[12]

Kornblith embraces a methodologically naturalistic approach to investigating the (purported) natural kind of knowledge. On this approach, the most compelling way to support the existence of a natural kind is to show that there is a branch of scientific inquiry that appeals to that kind in its causal-explanatory theorizing. In *Knowledge and Its Place in Nature*, Kornblith focuses specifically on "cognitive ethology," where this is understood as the scientific study of cognition in non-human animals. Kornblith offers a substantial case that ascriptions of knowledge to non-human animals do significant explanatory work in cognitive ethology, which cannot be done by (e.g.) mere belief-ascription, and which cannot be debunked as uncareful anthropomorphism by the relevant scientists.[13]

Kornblith takes this evidence from cognitive ethology to support a specific substantive understanding of knowledge: roughly, that knowledge is reliably formed true belief.[14] To a first approximation, this is because, when explaining animal behavior, we need to distinguish mere true beliefs from those which are a result of the animal's "attunement" to its environment: that is, true beliefs that are reliably formed.[15]

Kornblith offers an account of the normativity of epistemic norms that rests on his claim that people have a universal interest in being able to make accurate assessments regardless of whatever values they might happen to have.[16] In turn, he argues that this interest provides instrumental reason for people to care about knowledge. For Kornblith, that captures and explains at least much of the "normativity" that matters in theorizing about knowledge.[17]

Finally, Kornblith appears to accept a "causal or historical" metasemantics for the word 'knowledge'.[18] On this sort of account, the content of this word is fixed, not by our intuitions or judgments about it, but by the word's robust causal connections to the natural kind mentioned above. One worry about a theory that focuses on a relation that plays

a central role in cognitive ethology is that it in some sense "changes the topic" from the subject that epistemologists are ordinarily focused on. We take it that this metasemantics could potentially help to defuse this worry: on this view, ethologists are investigating the very relation—knowledge—picked out by the word 'knowledge', the use of which is central to contemporary epistemic discussion.

With these six commitments in hand, we can see how Kornblith's work promises to achieve the two goals we have identified for epistemic inquiry. Very briefly, his identification of knowledge with a specific (alleged) natural kind—roughly, reliably formed true belief—aims to capture the explanatory importance of epistemic entities. And his idea that epistemic norms are "universal hypothetical imperatives" aims to capture what gives epistemic norms their distinctive normativity.

Each of the six commitments we have sketched here is deeply controversial. However, we also take each of them to have some independent plausibility, and thus to be worth taking seriously. Moreover, they together constitute a unified and seemingly explanatorily powerful approach to epistemic theorizing. Our approach in what follows is largely to grant Kornblith (much of) this package of commitments, in order to argue that there are significant worries about Kornblith's approach, relative to the explanatory and normative goals that we take to be central to an attractive form of object-level epistemic inquiry.

4 Explanatory Importance

Suppose that we grant Kornblith that he has identified a natural kind; that (further granting his metasemantic commitments), 'knowledge' picks out that kind; and, hence, that (to semantically descend) we can thus identify knowledge with that kind. One might ask: given such an extraordinary level of success, what more could a naturalistically-minded object-level epistemologist want? In this section and the next, we argue that such an epistemologist might still be dissatisfied with how well Kornblith has addressed the aims of epistemic inquiry that we introduced earlier. We start, in this section, by focusing on explanatory aims. We raise two worries. The first arises because Kornblith focuses his inquiry on a folk epistemic category: namely, knowledge. The second arises because of Kornblith's limited engagement with the range of apparent explanatory roles that epistemic entities can play.

4.1 Explanatory Depth and Folk Categories

Recall Kornblith's slogan that "the subject matter of epistemology is knowledge itself, not our concept of knowledge."[19] In the preceding section, we emphasized how this slogan exemplifies a commitment to object-level inquiry. Taken literally, however, this slogan has another possible implication: that the study of knowledge exhausts epistemology. We take this idea to be a non-starter. This is because epistemology plausibly includes the study of a range of things, including (at least plausibly) epistemic justification, evidence, epistemic rationality, and epistemic praise and blame. This observation allows us to ask an important

methodological question: how central is the study of *knowledge* to object-level epistemology? Indeed, we might ask: is it really that important at all? This question might sound absurd. But, we argue, it can be motivated by appeal to some of Kornblith's own commitments.

In explaining why he sees no significant role in epistemology for the explicit investigation of epistemic concepts, Kornblith asks a pointed rhetorical question: "Why should our folk epistemic notions be of any more interest to epistemologists than our folk chemical notions are to chemists?"[20] Against the standard posed by this question, Kornblith's own attempt to detach himself from folk epistemology is crucially incomplete. This is because, while he does not attempt to offer a systematic account of the *meaning* of any folk epistemic terms, he targets his inquiry as a kind that he takes to be the referent of a folk epistemic term: 'knowledge'.

We take this to suggest a significant worry, even if we grant that knowledge is a natural kind. To see this worry, we can develop Kornblith's analogy to chemistry. Suppose (for the sake of argument) we grant that the folk term 'water' picks out a chemical natural kind, H_2O.[21] Arguably, water often features centrally in both contemporary "folk chemistry" and historical, proto-scientific chemistry. It would be a mistake to infer from these facts that this natural kind is of central theoretical importance to chemistry. Rather, it turns out to be somewhat theoretically marginal in chemistry, compared to (for example) elements like hydrogen and oxygen.[22] We can now apply the analogy back to epistemology: the term 'knowledge' is a bit of "folk" vocabulary. Indeed, it is arguably the folk epistemic word *par excellence*. As a folk term, even if it turns out to pick out a natural kind, it might be one that is marginal to the explanatory structure of the epistemic.[23]

This possibility suggests that there is a "folk representational-level" residue in Kornblith's approach to epistemology. That is, he seems to assume that if the folk word 'knowledge' picks out a natural kind, then that kind is of central epistemic importance. This assumption seems needed to explain the effort Kornblith dedicates to defending the idea that the relation he is focused on is *knowledge*. But, as the analogy to chemistry shows, naturalistically-minded philosophers should be suspicious of this easy inference from centrality within folk thought and talk to theoretical importance.

Consider another way to press this point. As we have noted, Kornblith argues for his account of knowledge as a natural kind largely through discussion of cognitive ethology. Suppose it turns out that parts of scientific inquiry other than cognitive ethology also involve scientists attributing "knowledge" to certain entities (whether humans, non-human animals, robots, or other entities). And suppose that these different parts of scientific practice track natural kinds that are different from the natural kind that figures in our best accounts of cognitive ethology. To appreciate some of why this possibility is worth taking seriously, consider some of the myriad ways in which researchers in artificial intelligence talk about neural networks or robots "knowing" certain things (without it being clear that these things have "beliefs" in a more familiar sense that we attribute to both human and non-human animals), or how many social scientists (including, for example, many historians and psychologists) appear to use 'knowledge' in non-factive ways (on which people might "know" things that are false, given the social/psychological role of the relevant beliefs).[24] These different uses of 'knowledge' might all turn out to pick out different natural kinds. If so, one question is which of these kinds (if any) is really what

our folk term 'knowledge' refers to. But the more pressing question for the epistemologist who embraces our goals is this: which of these kinds (if any) are explanatorily important? And, if multiple ones are, which ones are *more* explanatorily powerful or central? Even if we grant that Kornblith is correct about the referent of 'knowledge', this does not settle these further questions about explanatory importance.

We want to be clear that here, as is typical throughout this chapter, we are posing a *methodological* challenge. We do not do so with a specific foil for Kornblith's substantive epistemic views in hand. Rather, we aim to emphasize that Kornblith appears to structure his inquiry in a way that could prevent him from achieving the two goals for object-level epistemology that we have advertised. Part of the reason for this, we have suggested, is that he appears to be operating with an assumption that is in tension with a natural and appealing extension of his own naturalistic commitments: the assumption that the appropriate focus for our central epistemic investigations is the referent of a folk term.

4.2 Explanatory Power and the Regulative Roles of the Epistemic

One salient difference between epistemology and chemistry is that we might take epistemology to have a distinctively *regulative* character for human life. To illustrate, consider that epistemic entities seemingly (both do and should) play roles in regulating the following things (among many others): social practices of inquiry;[25] individual belief formation and maintenance;[26] assertion and action;[27] and practices of praise and blame.[28] Later, we consider normative issues about which things should play such "regulative" roles. Here, we focus on the fact that, if some epistemic entities in fact do play these sorts of regulative roles, this is an important descriptive fact about them, which will matter for what kind of explanatory importance those entities have. To help see this point, consider that for an entity to play a regulative role in the relevant descriptive sense is plausibly for it to play a nomically significant explanatory role. Because Kornblith understands natural kinds as the sorts of entities capable of standing in nomically stable relations, this would be evidence that the relevant epistemic entities are natural kinds (on Kornblith's understanding of "natural kinds").

A central element of Chapter 3 of *Knowledge and Its Place in Nature* is Kornblith's argument against an important idea that runs through much of contemporary epistemology, which is influentially defended by Ernest Sosa, among others.[29] This idea begins by suggesting that we can distinguish between a sort of "knowledge" that can be possessed by non-human animals, and a richer sort of "knowledge" which requires capacities characteristically available to adult humans, but not to most non-human animals (or to younger human children). It then suggests that epistemic inquiry should centrally focus on the latter sort of "knowledge." Sosa characterizes this richer sort of "knowledge" as distinctively "reflective."[30] But one could take seriously the idea of a distinctively human sort of "knowledge" even if (following Kornblith) one is suspicious of the idea that it centrally involves reflection. In response to this kind of idea, Kornblith says that we may, if we wish, draw a distinction between "Vermont

knowledge" and "mere out of state Knowledge" (i.e. the knowledge had by people living outside of Vermont). But, he argues, this distinction (somewhat obviously) "marks no significant difference," relative to the concerns of philosophers doing epistemology.[31]

It is striking that when Kornblith seeks to criticize the idea of a richer sort of "knowledge" characteristically available to human adults, the core of his criticism involves casting doubt on a distinctive *normative* role for such a relation. This is especially true of his discussion of the practice of giving and asking for reasons. Kornblith wants to emphasize that this practice is good only in certain conditions, and otherwise counterproductive.[32]

Consider instead the following thought, which is purely about the explanatory importance of *descriptive* matters, which is our focus in this section. Human adults characteristically have more sophisticated cognitive and communicative capacities (at least on certain salient dimensions of those capacities) than many non-human animals. Those capacities might allow human adults to have a knowledge-like state that is richer than mere reliably formed belief, and which could be part of a "homeostatic cluster" that is explanatorily potent. Adding credibility to this possibility is the fact that (something like) knowledge appears to play a regulative role in the series of distinctively human practices mentioned at the beginning of this subsection. (We arguably do not see non-human animals engaged in full-blown assertion or blame, for example.) We can contrast this with Kornblith's example of "Vermont Knowledge": this example is compelling in part because Vermonters *do not* have strikingly more sophisticated capacities and practices than other humans. Together, the observations laid out in this paragraph make credible the hypothesis that a distinctively human "knowledge-like" relation may play a range of important explanatory roles not played by "mere animal knowledge." If this hypothesis is correct, it would entail that the former relation was a genuine epistemic natural kind. And this in turn would entail that a naturalistically-minded epistemologist would have reason to investigate that kind.

Notice again that this observation is wholly compatible with the commitment that we are granting Kornblith: namely, that the word 'knowledge' picks out what some call "mere animal knowledge." Our point is that, from the perspective of an epistemologist with the kind of object-level focus we are advocating for, this fact about our language hardly settles the issues at hand.

This section has made two key points. First, from a Kornblithian naturalistic methodological perspective, there is no particular reason to think that our folk term 'knowledge' will pick out an especially central epistemic category. Second, the combination of distinctive human capacities and distinctive epistemic practices that humans engage in provides reason to think that there may be epistemic kinds that are central to explaining those practices, and that extend beyond knowledge as Kornblith theorizes it. Together, these arguments raise a striking possibility: Kornblith's six commitments could all be correct, insofar as they concern knowledge, and yet Kornblith could fail to provide a theory of many of the *most* explanatory important things in epistemology, and could instead provide an account of something relatively explanatorily marginal.

5 The Normative Significance of the Epistemic

We now consider how well Kornblith's project addresses the second goal for object-level epistemic inquiry that we have advertised: accommodating and illuminating the normative character of the epistemic. Our central concern is that Kornblith is not sensitive enough to the range of normative explananda to be addressed. We focus on two issues: the rich variety of normative roles seemingly played by epistemic entities, and the apparent normative "authoritativeness" of epistemic norms. We conclude the section by considering, as a case study, the methodological significance of "Gettier Cases."

5.1 The Apparent Diversity of Epistemic Normativity

In Section 4.2, we noted the descriptive point that epistemic entities (or our representations of them) seemingly play central roles in the regulation of a wide range of phenomena. For example, they appear relevant to the regulation of individual belief, social practices of inquiry, assertion, action, praise, and blame, among other things. In this section, we emphasize the diverse *normative* relations that epistemic entities seemingly bear to these phenomena. Consider some initially plausible examples: a person who has failed to apportion their belief to the evidence has made a distinctive kind of *mistake*; assertion is (in some sense) *correct* when the asserted content is known; evidence constitutes *normative reasons* to believe; knowledge is more *valuable* than true belief; intellectual humility and creativity are epistemic *virtues*; and belief can be held epistemically *permissibly,* or *responsibly*.

Here is a natural methodological thought: an account of epistemic normativity should explain a range of actual normative phenomena tied to the epistemic, such as those listed above. Of course, some of the apparent phenomena, listed above, may be in some sense illusory, and thus may not need to be accounted for. But if at least many of them are on track, we take the broad methodological upshot to be this: we should be aiming to give a satisfying explanation of many of these phenomena, and (ideally) a satisfying debunking of the apparent phenomena that we cannot explain.

So, we can ask: is Kornblith's account well-placed to offer such an explanation? It is somewhat hard to tell, because his core discussion of normativity is largely silent on this range of phenomena. We are inclined to uncertainty, tinged with pessimism. One reason for this stance is that notions of reasons, rationality, and the deontic categories have featured centrally in arguments for alternatives to reliabilist approaches to epistemology.[33] This suggests that Kornblith's reliabilism may be a poor framework for accounting for the full normative texture of the epistemic.

Suppose that Kornblith's account is incapable of explaining some of the relevant normative phenomena. This could be evidence that his account fails qua account of knowledge. But we want to emphasize a different possibility. Even if we suppose that Kornblith really has given us a correct account of the normativity of *knowledge*, this account might simply show that this folk category is not the most theoretically crucial or central epistemic category. Perhaps, indeed, the most important normative statuses are borne by epistemic categories for which we lack folk terms.

Of course, epistemologists widely believe that *knowledge* is normatively significant for the phenomena we glossed above. But this could be explained by the facts that (1) there is an epistemic entity that does bear these sorts of significance, and (2) these epistemologists are insufficiently attuned to the specific entity that causally regulates our use of 'knowledge', and are incorrectly using 'knowledge' in theorizing that is in fact sensitive to some other epistemic entity.[34] It seems to us that, merely in virtue of correctly characterizing the normative significance of the folk epistemic kind (as we are assuming), Kornblith could be falling well short of what we might hope a naturalistic theory of epistemic normativity might achieve: that is, to provide a theory that accounts in a robust way for the range of normative phenomena to be explained.

5.2 Authoritativeness and Epistemic Normativity

Kornblith's defender might reasonably protest at this point: even if there is a range of such normative phenomena, provided that Kornblith has accounted for the sort of epistemic normativity that *really matters* for how we live our lives, then his account may be incomplete, but it would still be a profound accomplishment. We agree, and in this section, we take up this idea, by focusing on the broad idea of norms that "really matter."

To begin, consider an intuitive contrast between ethical norms, on the one hand, and the rules of a silly party game, on the other. In one sense, these are both "normative": for example, perhaps you can break the rules of the party game by failing to refer to yourself in the third person, or by removing your hand from the top of your head. But in another, natural sense, there seems to be a contrast here: the ethical norms seem to normatively "matter" in a way that the party game rules do not. In some of our recent work, we have proposed a way to regiment this contrast using a distinction between "merely generic" and "authoritative" normativity. Here's the basic distinction. Take "generic" normativity to be the kind of normativity that *any* set or rules, standards, or norms have, from rules of board games to moral standards. In contrast, "authoritative" normativity is meant to be a kind of normativity that people (at least sometimes) appeal to in making claims about what one "really and truly should do," or what is valuable *simpliciter*. The idea is that, in some of the latter cases, people seem to be reaching for a kind of normativity that is maximally "normatively important." Put roughly, we might think this is the sort of normativity that, were it to exist, would be possessed by a norm that is suitable to normatively settle perceived conflicts between all other sets of norms (e.g., a perceived conflict between what the legal and moral norms prescribe).[35]

With this distinction in hand, note that many epistemologists take (at least some) epistemic entities to have some kind of deep tie to authoritative normativity. This tie might be because the epistemic is just a subset of the authoritatively normative. Or it might be because the relevant epistemic facts entail authoritatively normative ones, or because they play a major role (of the right kind) in helping determine them. The point is just that the relevant epistemic entities seem (at least to many) to have some kind of tie to authoritative normativity, which contrasts with the party game rules. Arguably, this (perceived) connection to authoritative normativity motivates many (though not all) epistemologists to study

epistemic norms, in a way that parallels the motivations for why many philosophers study (for example) moral standards and principles of distributive justice.[36]

With this in mind, we can ask: has Kornblith given us a plausible account of the epistemic norms' connection to authoritative normativity? Recall the core of Kornblith's account of the normativity of epistemic norms. On his account, epistemic norms enable accurate assessments. Because of this, they have a kind of "hypothetical" normativity: roughly, *if* one wants to be able to make accurate assessments, then one should follow these norms. Kornblith suggests that everyone has an interest in being able to make accurate assessments, regardless of whatever other values they might happen to have. And this interest in turn entails that, while hypothetical, epistemic norms apply universally to human agents.

The crucial question is whether this sort of "hypotheticalist" account of epistemic normativity suffices to explain the seeming authoritativeness of the epistemic norms.[37] One general basis for suspicion is that many metaethicists find hypotheticalist accounts of moral or practical normativity unpromising. We can give a brief overview of some of the discomfort here, by considering a parallel case. Humans are social creatures, such that securing the cooperation of other humans is very often (indeed, perhaps almost always) helpful for promoting one's interests. In light of this, successfully exercising the capacity to manipulate others into doing what one wants will reliably promote one's interests. But it is implausible to conclude on this basis that we have vindicated the authoritative normativity of the norms that promote successful manipulation of other humans. Indeed, arguably, the latter norms have no deep tie to authoritative normativity.

It is possible that Kornblith's style of account can be developed with more nuance to address this familiar sort of challenge.[38] But consider what we take to be a live possibility (and, indeed, perhaps the *more* likely possibility): that universal hypotheticalism cannot be developed into a plausible account of authoritative normativity. Then it is possible that Kornblith could give the correct account of the normativity of *knowledge*, but that knowledge is not the locus of what normatively matters most in epistemology. To press this point, suppose that there are some epistemic entities other than knowledge (possibly ones for which we lack folk terms) that are authoritatively normative. If this were the case, Kornblith's having given a correct account of *knowledge* would be a somewhat marginal achievement from the perspective of an object-level epistemologist concerned with theorizing about the epistemic norms that really matter.

5.3 Gettier Cases

In many contexts, it might well be that an epistemologist adopting our proposed approach to object-level inquiry within epistemology—as opposed to others that are centered more on issues about the reference of folk concepts and words—might not change much (if anything) about the work she does. After all, it might be that different approaches converge on methodological recommendations in many contexts, or that the actual practices of epistemologists turn out not to be that impacted by their explicit methodological ideas.

However, we think that adopting our approach (as opposed to others, which are more tethered to facts about our existing thought and talk) might well matter. To illustrate, we turn to how our approach highlights important (and non-standard) ways of thinking about recent discussions in epistemology. We focus on a particular case study: namely, the kinds of cases that stem from Edmund Gettier's famous discussion of knowledge in "Is Justified True Belief Knowledge?"[39]

In that paper, Gettier introduced a series of cases where someone appears to have justified true belief about a proposition, but (allegedly) fails to know it. In the wake of Gettier's discussion, epistemologists have identified an extraordinarily rich collection of alleged counterexamples to various theories of knowledge. Many working in epistemology take these to constitute a *family* of cases, which are in some sense "of a kind" with Gettier's original cases.[40] Call this family *Gettier Cases*. Discussion of these cases structured much of epistemology in the late twentieth century, and they continue to play an important role in the field today. Given this prominence, it's striking how little attention they get in *Knowledge and Its Place in Nature*. Indeed, in the book, Kornblith only explicitly mentions such cases in his initial discussion about the potential significance of the appeal to intuitions for "naturalistic" approaches to epistemology (which he sees as casting doubt on the importance that Gettier Cases are often given in philosophical practice).[41]

This lack of attention to Gettier Cases is surprising because the sort of straightforward reliabilist theory of knowledge that Kornblith appears to endorse in *Knowledge and Its Place in Nature* is surely subject to (apparent) counterexamples from within the class of Gettier Cases. Which specific counterexamples those are depends on how Kornblith spells out further details of his account of knowledge (including, crucially, details about his understanding of "reliability") that go beyond what we have space to discuss in this chapter. The core point is just that *many* examples in this literature were developed with the aim of undermining views of knowledge as reliably formed true belief in the ballpark of Kornblith's, including views that add in more complicated extra conditions involving "justification" or further developing the kind of reliability involved. We conjecture that Kornblith's disinterest could be explained as follows: he takes himself to have an independently robust case for his account of knowledge, and in light of this, he takes himself to have reason to hold that judgments about Gettier Cases are misleading, and that they play no significant evidential role.

Here, we continue our approach of granting that Kornblith has the correct theory of knowledge, in order to spell out what we take to be an important alternative role for Gettier Cases, given that assumption. Widespread Gettier Case judgments might be responsive to something (or things) of deep concern to many working in epistemology, *not* because they track the referent of the folk term 'knowledge', but rather because they help us to better understand something in the vicinity that is explanatorily or normatively important. (As usual, we do not assume that they always or even often do this, only that we should be attentive to the possibility that they do.)[42]

For someone pursuing the kind of object-level epistemic inquiry we are advocating for, the assumption that the word 'knowledge' fails to pick out a Gettier-proof relation should do little to undermine the interest of inquiry into such a relation. Instead, it potentially

suggests that *knowledge* is less central to the epistemic domain than we might initially have assumed. For example, perhaps a certain Gettier-proof relation plays an important role in explaining or rationalizing our practices of assertion, or in certain psychological explanations.[43] Or perhaps a certain Gettier-proof relation is (in some sense) *more valuable to have* than mere reliably formed belief.[44]

Gettier Cases can be methodologically valuable for object-level epistemic inquiry even if there is not a single, fully "Gettier-proof" condition. Consider the many proposals that philosophers have put forward about what knowledge is that seek to avoid being counter-exampled by Gettier Cases. In particular, focus on the history of proposals that seek to accomplish this task by adding extra conditions to a "justified true belief" model of knowledge, which was, for a long time, the main kind of response to such cases.[45] These proposals have helped to develop our understanding of a wide range of (purported) epistemic relations, which might well have significant normative importance. Consider the following examples: "safety" (roughly: for an agent's belief to be safe is for it to be such that the agent could not easily have falsely believed it) and "sensitivity" (roughly: for an agent's belief to be sensitive, the agent must fail to have the belief in the "nearest" worlds where the content is false). Thus, even if inquiry drawing on Gettier Cases has failed to help us to identify the correct theory of knowledge, and even if it has failed to identify a single kind that is quite generally "Gettier-proof," it might have helped us to better understand a range of normatively important epistemic kinds. Notice that the history of philosophical reflection on Gettier Cases might, in a parallel way, also have helped us better understand descriptive issues that matter for identifying and understanding epistemic entities that are explanatory important, whether they are also normatively important or not. Consider here, for example, attention to the "basing" relation (which is, roughly, the relation between a belief and the considerations for which the agent actually holds the belief).

Let us take stock. In this section, we have cast doubt on whether Kornblith has adequately theorized the normative significance of the epistemic. In doing so, we have again granted that he has correctly theorized knowledge. We have emphasized that this assumption is compatible with his failing to have adequately theorized the range of central normative phenomena connected to the epistemic, the authoritativeness of the epistemic, or the sorts of normative roles connected to the "Gettier"-style judgments that have played an important role in shaping work in epistemology over the last 70 years.

6 Representation and Object-Level-Focused Epistemology

In this section, we step back from our specific critiques of Kornblith to think more carefully about the relationship between representation-level inquiry, and the kind of object-level epistemic inquiry we are advertising. Our discussion has two parts. The first considers a potential objection to our critique of Kornblith, based on the idea that we are insufficiently attentive to preserving the *topic* of epistemic inquiry. The second explores what we take to be a central role for *normative* representation-level inquiry, given object-level theoretical goals.

6.1 Changing the Subject?

Frank Jackson provides a particularly clear version of a classic motivation for the object-level inquirer to engage in (at least some amount of) representation-level inquiry. For Jackson, inquiry into our thought and talk can help us to better identify and clarify the shared subject matter of our inquiry, and avoid accidentally changing the subject.[46] Jackson combines this idea with a specific view about what determines the content of our thought and talk. Put roughly, on Jackson's account, dispositions of individuals to apply (or not apply) terms play a key part in a holistic theory of content-determination. For Jackson, this means that (again, put roughly) a certain form of "conceptual analysis" that focuses on tracking those dispositions will be epistemically helpful for our understanding the contents of our concepts, and what we mean by our words. As we read Kornblith, he joins Jackson in holding that continuity with the relevant epistemic topics (e.g., the topic of "knowledge") can be secured by continuity of semantic content. However, as we have seen, Kornblith favors a kind of causal/historical metasemantic account, on which "subject matter is defined by way of connections with real kinds in the world," rather than by our beliefs or intuitions about "defining" features.[47]

The assumption about "topic continuity" that Jackson and Kornblith appear to share might seem to suggest a powerful objection to the type of critique that we have been offering of Kornblith. By seeking to decouple object-level epistemic inquiry not only from questions about our pre-theoretic "intuitions" about cases, but also from questions about reference, we seemingly risk "changing the subject" from what epistemologists currently study. In particular, if we are relentlessly focused on object-level issues about the explanatory and normative importance of various entities, and we do not stay tethered to facts about what people mean by terms like 'knowledge', 'justification', or 'evidence', could not we just end up talking about a different subject matter entirely? We offer a three-part response, which focuses on goals of inquiry, pre-theoretical judgments about topic, and our own theory of "topic continuity."

As an object-level inquirer, perhaps the most compelling reason to care about "changing the subject" from A to B is because you thought topic A is the one that you really should be thinking and talking about in the context at hand. Compare: it would arguably be absurd to complain about relativity theory in physics "changing the subject" from prior ways of thinking about mass or motion, if one is convinced that the concepts deployed in general relativity are ones we need to understand the physical structure of reality. We understand the possibilities that we have been offering in the preceding two sections in this way: whether or not ordinary people or epistemologists *have* been thinking about the most explanatorily and normatively important epistemic entities, we suggest that it would be good if they did. We take this sort of reply to be compelling, and especially so to someone with Kornblith's broad methodological orientation.

One might reply that merely focusing on explanatory and normative importance is not enough. For example, wildebeests and cacti are surely legitimate targets for inquiry aimed at explaining descriptive issues, and distributive justice and moral obligation are surely legitimate targets for normative inquiry. But that does not thereby make any of these topics part of *epistemology*. We begin by offering an intuitive, non-theoretical reply. We are not

proposing that epistemologists investigate wildebeests, or just do moral and political philosophy instead. Rather, we are suggesting that they focus on entities that are "in the vicinity" of what we currently refer to by the relevant epistemic vocabulary (e.g., 'knowledge' and 'justification'). And our conjectures are motivated by taking seriously apparent explananda and normative features of the epistemic that show up robustly in epistemic practice. For example, given the centrality of Gettier Cases to recent philosophical work within epistemology, it would be bizarre (in most philosophical contexts) to accuse an epistemologist of "changing the subject" for seeking to understand the relations sensitive to such cases.

We can bolster these observations by appealing to our own work on topic continuity. Our account stems from looking at the kinds of concerns that animate philosophical objections about whether given theories "change the topic", whether in epistemology, metaphysics, ethics, or elsewhere. These objections often invoke three sorts of central issues about thought and talk. First, there are issues about the *representational* roles a given part of thought and talk has. For example: what is the representational content of the word 'knowledge', and what things does it refer to across possible worlds, in light of that content? Second, there are issues about the *inferential* roles a given part of thought and talk has. For example: what kinds of inferences do people make when they think that someone else knows something, or when they hear someone else claim that someone else knows something? Third, there are issues about the sort of *communicative* roles a given part of thought and talk has. For example: when people use the term 'knowledge', what kind of direct and indirect ways of talking about relevant subject matters do people have? And how do their ways of communicating help them preserve "substantive" as opposed to "merely verbal" disputes?

In recent work, we have drawn on this pattern to develop an account of "topic continuity." Put roughly, on our account, when philosophers talk about "topic continuity," we think they should be tracking degrees of continuity in thought and talk along the three dimensions we have just briefly discussed: representational, inferential, and communicative. In turn, we think how much each of these different dimensions should matter relative to each other, and what amount of continuity one needs along a given dimension to reach a threshold for "preserving topic," will vary from context to context, based on the interests of the participants in that context.[48] If our account of topic continuity is on the right track, then, in certain contexts, philosophers can stay "on topic" when talking about "knowledge" (for example) even if they are no longer talking about the relation that the term 'knowledge' previously referred to. This result also holds for many other recent discussions of "topic continuity."[49]

It is also important to emphasize that inquiry is often concerned with a variety of topics. For example, we take it that many epistemologists aim in a given work to better understand both the topic "knowledge" and the topic "normatively important epistemic entities." For example, we think this is true of Kornblith in *Knowledge and Its Place in Nature*. We thus take it to be plausible that in many contexts, at least one of the topics we are focusing on here is already at issue in epistemic inquiry that is focused on knowledge.

We can turn briefly from defense to offense here. As we have emphasized, our account treats topic continuity as a matter of continuity of referential, inferential, and communicative functions. One upshot of our account is that even if Kornblith has correctly identified the

referent of 'knowledge', his account could fail to preserve the topic of "knowledge," in contexts where participants are especially interested in aspects of inferential and/or communicative functions associated with uses of 'knowledge' (either by ordinary people, or specialists of some kind) that are not secured by his account.[50] For reasons we explain in what follows, we do not take this to be a significant objection to his account, given a certain, plausible way of thinking about the aims of epistemic inquiry we have been discussing. But we do take it to further reinforce the dialectical awkwardness of objecting to our critique by appealing to topic.

Finally, we want to emphasize the modest role we see considerations of topic continuity playing within the sort of object-level inquiry that we are advocating for here. Recall our guiding idea (put at a high level of abstraction): that object-level epistemic inquiry should focus on parts of reality that have certain valuable properties, tied to the central goals of epistemic inquiry. Let us grant that being part of a certain topic might be one such valuable property for at least some epistemic inquiry (at least in certain contexts). This assumption leaves open the questions of *how* valuable this property is, relative to other properties, such as being particularly explanatorily or normatively important in some way. It might well be that, if we value these other things more, then, as our object-level inquiry in epistemology unfolds, we might move further and further away from the specific epistemic topics we are currently focused on within epistemology. And that might be entirely okay, given what we discover about what's really valuable to inquire into. This point applies to the topic of "knowledge," but also more broadly to the very topic of "the epistemic," however that is understood.[51] That is, it is conceivable that epistemic inquiry might virtuously evolve into a form of inquiry that leaves the current topic of "the epistemic" behind.

6.2 Constructive Inquiry into Representation for the Object-Level Epistemologist

We now want to argue for the relevance to the object-level epistemologist of *normative* inquiry into epistemic representations, of a sort that arises in "conceptual ethics" and "conceptual engineering."

To begin, we briefly explain what we take conceptual ethics and conceptual engineering to involve. On our understanding, conceptual ethics (put roughly) can be glossed as normative and evaluative inquiry into certain kinds of questions about thought and talk, such as ones about which concepts we should use, and why, or what it would be good for certain words to mean, and why.[52] Now turn to "conceptual engineering." Put roughly, the idea is this. On a given topic, a conceptual ethicist might conclude that the relevant current words and concepts (and our use of them) should be retained as they are. But in other cases, they might conclude that the relevant current concepts (or words, etc.) should be reformed or replaced in some way. Paradigmatic conceptual engineering projects begin with such "reforming" or "revolutionary" conclusions in conceptual ethics. They then involve work on crafting a concept (or word, etc.) that reflects the desired changes (what we can call "conceptual innovation"), and attempts to get relevant people to start using that proposed concept or word in the desired contexts (what we can call "conceptual implementation").[53]

With these understandings of conceptual ethics and conceptual engineering in hand, return to the kind of object-level inquiry in epistemology we have advocated for. So far, we have emphasized a natural desideratum for object-level inquirers (whether in epistemology or elsewhere): they should inquire into the entities that substantively matter, even if there is no folk term for them. For example, we have argued that if there is a "Gettier-proof" relation that is an epistemically important kind (whether because of its explanatory importance, normative importance, or both), then epistemologists should inquire into that relation. It is natural to think that we should also have *representations* that enable us to think and talk about this kind in fruitful ways: for example, by directly thinking and talking about it. That's a point in conceptual ethics, which, depending on where things stand with our existing epistemic thought and talk, might naturally lead to work in conceptual engineering to develop such representations.

Consider three further sorts of considerations that bear on questions about which concepts and words to use in epistemology. First, we are considering concepts and words to be used by *us*. We are finite creatures with various cognitive limitations and biases. Ideally, the relevant concepts and words should fit us, in that they allow creatures with our capacities and limitations to make progress in our inquiries and projects. Second, in Sections 4 and 5, we emphasized the range of explanatory and normative roles apparently played by epistemic entities. Notably, many of these roles appear to be played in part via epistemic thought and talk: for example, we are sensitive to epistemic norms in part by being sensitive to representations of those norms. So, some of the relevant epistemic representations—namely, the ones that are theoretically valuable in light of certain connections to these explanatory and normative roles—need to be suitable to play these roles. Finally, philosophers working on these issues face a choice that is familiar from a range of debates about conceptual ethics in other areas: when we identify a concept (e.g.) that is valuable to use in some of the above ways, do we want to introduce a new word into the lexicon to talk about it, or do we want to pair this new content with an existing term (as in practices of "explication" or semantic "precisification")?[54]

The upshot of all of this is that we think that even an inquirer who is resolutely ultimately focused on object-level issues should reject a literal reading of Kornblith's slogan that "epistemology should not concern itself with our concept of knowledge."[55] This is not because the ultimate aim of the kind of epistemic inquiry we have been advocated for concerns "representational-level" issues. It is certainly not because its ultimate aim is to accurately describe a folk concept. Rather, it is because normative and evaluative inquiry into our actual epistemic thought and talk—as well as into possible alternatives to that thought and talk—can play crucial roles in object-level epistemic inquiry that ultimately aims to illuminate the most (explanatorily and normatively) important epistemic entities.

7 Conclusion

In this chapter, we have advertised a certain form of object-level inquiry in epistemology. We have introduced two goals, which we claim reflect central explanatory and normative interests in epistemology. We then argued that, by taking these goals seriously, we can

identify a series of concerns about the adequacy of the epistemic theory Kornblith offers in *Knowledge and Its Place in Nature*. We have explained how, even if we grant that Kornblith has successfully characterized the nature of knowledge, his theory could still be disappointing relative to the goals we have identified: namely, it could fail to theorize the most explanatorily and normatively important parts of "epistemic reality." We explained why, given Kornblith's account, knowledge could turn out to be a theoretically marginal natural kind in epistemology, much as water is in chemistry. We also explained why it could turn out to be normatively marginal. Perhaps more plausibly, even if it is not explanatorily or normatively marginal, it could turn out that knowledge is less important on one (or both) of these dimensions than other epistemic kinds.

Our discussion suggests a consistent diagnosis for why Kornblith's work on knowledge might fail to meet the goals of epistemic inquiry that we have been discussing. In short, at the core of the explanation for the possibility that Kornblith's work in epistemology might end up achieving superficial success yet suffering deeper failure is the fact that, despite its seemingly radical object-level character, Kornblith's inquiry is not as resolutely "object-focused" as it could be. Rather, it is consistently structured, without significant interrogation of this choice, around the relation referred to by a central folk term: namely, 'knowledge'.

We submit that the possibilities we have sketched suggest a methodological upshot: the naturalistically-minded, object-level epistemologist has strong reasons to investigate whether epistemic entities exist that can play the range of explanatory and normative roles we have identified here *better* than the entities picked out by our ordinary epistemic thought and talk. If the former sorts of entities exist, epistemologists should seek to answer a range of explanatory and extensional questions about them. This work might include, for example, investigating their natures, grounds, and supervenience bases.[56] We have also argued that, once we notice that the object-level epistemologist cannot simply take our folk epistemic terms for granted, there is a natural motivation for them to engage in work in conceptual ethics and conceptual engineering, to identify the representational-level resources that best allow us to investigate, understand, and make use of the epistemic entities that object-level inquiry reveals.

In what follows, we step back from our critical discussion, to locate that discussion within our own broader methodological approach to epistemology. In our view, much philosophical inquiry is (often implicitly) structured by the assumption that the meanings of our folk terms, the substantial content of our "intuitions," what is explanatorily important in the vicinity, and what is normatively important in the vicinity, all "line up." Call this the "alignment assumption." In much of his work, Kornblith emphasizes a possible disconnect between ordinary "intuitions" and the character of the explanatorily and normatively important kinds in epistemology. In this chapter, we have sought to emphasize a similar possible disconnect between the referents of our folk epistemic terms and the relevantly important kinds.

We can step back from these specific arguments to notice that, in general, we should be alive to the possibility that the above "alignment assumption" is mistaken. If it is, there would then be rich theoretical questions about how these different categories—meanings, intuitions, explanatory importance, and normative importance—relate to each other.

If they come apart, there will be interesting theoretical choices concerning what among these categories is worth investigating, and why. Especially tough questions arise when there are issues about which of these categories should play which sorts of roles in our lives. For example, consider the question of which of these categories to prioritize when evaluating which kind of epistemic entity should play a certain central role in regulating practices of epistemic praise and blame.

For simplicity, this chapter has focused on advertising a specific conception of object-level epistemic inquiry, which takes a general stand on the methodological questions just posed. Our own considered views are far more methodologically pluralist than this conception, along several dimensions. Thus, while we think that epistemic inquiry focused on "epistemic reality" is important, we do not think it's the only project in epistemology worth pursuing. For example, consider what we discuss in other work as the project of "metaepistemology." That is the project of explaining how actual epistemic thought and talk—and what (if anything) such thought and talk is distinctively about—fits into reality.[57] There is also the project of the conceptual ethics of epistemology, which we have just been discussing. In our view, epistemology focused on object-level issues (which can often be understood as a part of "normative epistemology"), metaepistemology, and the conceptual ethics of epistemology are all important projects.[58]

Moreover, on our view, these forms of inquiry can also inform each other in a variety of different, often subtle ways.[59] We have seen some examples in this chapter. For example, we have seen some of how the conceptual ethics of epistemology connects to the sort of object-level inquiry we discuss here, and how work on what we currently mean by our words can matter for issues about "topic continuity" and "changing the subject" that emerge from the literature on conceptual engineering. In other work, we explore further connections, including the ways that work in metaepistemology can inform normative, object-level inquiry in epistemology by impacting our assessment of what an apt epistemology for that later project might or might not be, which can then in turn matter for the viability of certain substantive, object-level views versus others.[60] As this example suggests, we think that the kind of object-level inquiry we advocate for in this chapter might well benefit in significant ways from—and also help inform—a range of other philosophical projects in epistemology that we think can also be well worth pursuing.

Our hope is that, if nothing else, our work in this chapter can help alert philosophers to this range of goals and projects, and encourage them to clarify which of these goals and projects—or alternatives we have not mentioned—they aim to contribute to, and why. Such clarification can then usefully inform what kinds of arguments they make, and what sort of evidence they look to in support of their theories.

Acknowledgments

Thanks to Chloe Banino, Christine Bratu, Ray Briggs, Herman Cappelen, Matti Eklund, Rachel Fraser, Sebastian Köhler, Niko Kolodny, Catherine Lapey, Maria Lasonen-Aarnio, Cara Lewis, Erika Milam, Philip Pettit, Dietmar von der Pfordten, Adrian Russian, Thomas Schmidt, Patrick Todd, and Marcus Willaschek for helpful feedback and discussion.

Earlier versions of this chapter were presented at Bielefeld University, Frankfurt School of Finance and Management, University of Hamburg, University of Hong Kong, Humboldt University of Berlin, University of Göttingen, and Uppsala University. Thanks to everyone who participated in those discussions. Thanks also to the editors of this volume, Joshua DiPaolo and Luis R.G. Oliveira.

Notes

1 We say "appear" because if certain views about epistemic thought and talk turn out to be correct (such as, for example, certain kinds of expressivist or pragmatist views), this appearance might be either defeated or at least significantly complicated. In this chapter, we largely set aside such possibilities. For further discussion, see McPherson and Plunkett (2021).

2 In this chapter, we use small caps (e.g., RATIONAL) for concepts, single quotation marks (e.g., 'rational') to strictly mention words, and double quotation marks (e.g., "rational") for a variety of tasks, including quoting other people, scare quotes, and mixes of use and mention.

3 See Jackson (1998) for sympathetic discussion and development of this basic idea. (It should be emphasized that, in developing this idea, Jackson ends up with a more subtle view of the targets of philosophical inquiry than the one we are glossing here).

4 (Kornblith 2002).

5 For a version of the first idea, see Haslanger (2000). And for a version of the second idea, see Rorty (1980).

6 Our discussion in this section draws on (McPherson and Plunkett (n.d.) Manuscript). In that paper, we discuss a range of different kinds of goals one might have for object-level inquiry in philosophy, how the goal one pursues impacts which objects one studies in philosophy, and how reflection on these issues can help us better understand and make progress on a range of philosophical debates (including ones about realism and pragmatism).

7 It is worth emphasizing that we are not assuming that guidance necessarily involves some sort of distinctively reflective activity. (This point is especially important in the context of discussing Kornblith, given Kornblith's critiques of the value of reflection (see, e.g., Kornblith [2012]). Here's an analogy: intensive collective inquiry into chess theory has enabled the best chess players to become much stronger than players in previous eras. But at the chessboard, the fruits of this inquiry will largely manifest in strong players' abilities to automatically "see" vulnerabilities etc., and not in reflective application of strategic principles.

8 (Kornblith 2002, 1).

9 For discussion, see Jackson (1998).

10 (Kornblith 2002, 1).

11 (Kornblith 2002, 61). Kornblith's account of natural kinds draws significantly on Boyd (1988/1997).

12 (Kornblith 2002, 62).

13 (Kornblith 2002, §2.5–2.6).

14 (Kornblith 2002, 63).

15 (Kornblith 2002, 58).

16 (Kornblith 2002, §5.5).

17 (Kornblith 2002, 161) suggests that it "remains to be seen" whether any sense can be made of aspects of epistemic normativity that his account alone cannot explain.

18 (Kornblith 2002, 11, 23).

19 (Kornblith 2002, 1).

20 (Kornblith 2002, 19).
21 For reasons to doubt this claim, see Weisberg (2006).
22 Note that this is consistent with water being a relatively important kind in other areas of the natural and social sciences. Our point here is just about chemistry.
23 To be clear: identifying the referent of 'knowledge' can be a central achievement within the project of understanding how our actual epistemic thought and talk works, which we take to be a central part of "metaepistemology." As we emphasize later in this chapter, we think this can also be a valuable project within philosophy, whether or not it turns out to be instrumentally helpful for the particular kind of object-level inquiry we are advocating here. For further discussion of metaepistemology, see McPherson and Plunkett (2021).
24 For a prominent example of this latter kind of non-factive use of "knowledge," see Gilovich (2001).
25 Here, consider especially Edward Craig's influential idea that designation as a knower plays a central role in decisions about whose testimony to trust. See Craig (1990).
26 Some philosophers have criticized "mainstream" epistemology for inadequacy on this front. For example, see Bishop and Trout (2004) and Ballantyne (2019). For discussion and defense of mainstream epistemology on this front, see Fantl (2023).
27 For the idea that knowledge is the norm of assertion, see Williamson (1996). For the idea that knowledge norms play a central role in rationalizing action, see Fantl and McGrath (2009).
28 See, e.g., Johnson King (2022).
29 (Kornblith 2002, ch. 3).
30 See, e.g., Sosa (2009).
31 (Kornblith 2002, 73).
32 (Kornblith 2002, §3.4).
33 See Pappas (2023, §8) for discussion of these themes.
34 In other work, we emphasize a different possibility, which connects to this idea that inquirers might be thinking about some other epistemic kind. This possibility is that inquirers could sometimes be implicitly switching between thoughts about different kinds without realizing it (McPherson and Plunkett 2020). We think that this danger is especially salient with terms that have both "folk" and "technical" uses, as many epistemic terms do.
35 For further discussion, see McPherson (2018) and McPherson and Plunkett (forthcoming).
36 Our claim that many in epistemology (either implicitly or explicitly) take there to be some important tie between epistemic entities and "authoritative" normativity does not mean that everyone working in the area does. For example, for a dissenting view, see Maguire and Woods (2020).
37 One worry about this account is that it is not clear that people's (alleged universal) interest in beneficial actions generates a strong instrumental reason to care about accuracy, per se, as opposed to, for example, caring about making "accurate enough" assessments. For connected discussion, see Nolfi (2020).
38 We take Mark Schroeder's version of hypotheticalism to show the most promising way forward on this issue. See Schroeder (2007).
39 (Gettier 1963).
40 To get a sense of the broad scope of such "Gettier Cases," and how they have mattered for work in epistemology, see Hetherington (2018).
41 (Kornblith 2002, ch. 1). For connected discussion, see Kornblith (2013), in which Kornblith downplays the role of "armchair methods" in epistemology, and argues that it is likely that work in the cognitive sciences is needed to achieve "real enlightenment."

42 For connected reflection on this kind of possibility, see Risberg (2023).
43 For example, Williamson's burglar case (Williamson 2000, 62), specifically appeals to one candidate Gettier-proofing feature in arguing for the superiority of knowledge over true belief in some psychological explanations.
44 Explaining how a Gettier-proof "knowledge" relation might be more valuable than "mere" justified true belief is part of what is sometimes called the "secondary value problem" in epistemology. See Pritchard (2007).
45 For an alternative kind of account, which departs from this tradition, and which sees knowledge as explanatorily primitive, see Williamson (2000).
46 (Jackson 1998, ch. 1).
47 (Kornblith 2002, 12).
48 See McPherson and Plunkett (2024b).
49 See, for example, the views offered in Cappelen (2018) and Thomasson (2020).
50 Indeed, in conversation, several philosophers have suggested to us that Kornblith is not talking about the thing they are interested in understanding when discussing "knowledge."
51 It is worth noting that there are a wide range of ways that philosophers think about what unifies properly "epistemic" considerations. For example, some identify this category via ostension, others via facts about constitutive standards governing belief, and others via connections to truth-centered considerations, of the right kind. Our basic methodological point here applies however one understands the boundaries of the category.
52 Our use of 'conceptual ethics' in this way draws from Burgess and Plunkett (2013a) and Burgess and Plunkett (2013b). Two notes on the label 'conceptual ethics' here, in light of ways that it might be misleading. First, as our discussion of issues about words underscores, conceptual ethics need not focus just on issues about concepts. It can be about a range of components of thought and talk (and which ones should be on that list, and how generally important they are relative to each other, is a theory-dependent matter, with significant disagreement between philosophers). Second, the term 'ethics' is meant here as a useful shorthand for "normative and evaluative inquiry." So it's not as if, for example, only broadly "practical" considerations, values, and norms matter here, of the kind that find their natural home in moral and political philosophy, such as the value of promoting happiness or being in accordance with the demands of distributive justice. Other things that are more tied to "theoretical" concerns can matter too: such as broadly "epistemic" goods about promoting successful inquiry, or broadly "metaphysical" ones such as carving reality at its joints. For further discussion, see Burgess and Plunkett (2013a, 2013b).
53 Our gloss of "conceptual engineering" follows the basic accounts given in Burgess and Plunkett (2020) and Cappelen and Plunkett (2020). Those two accounts differ in some details, but those details don't matter to us here.
54 For other work in epistemology (or other related normative domains, such as ethics) that advocates for the importance of (what we are calling) "conceptual ethics" and "conceptual engineering," which connects to our main line of discussion in interesting ways, see Haslanger (1999), Scharp (2020), and Eklund (2017). For further discussion of "conceptual ethics" readings of recent work in epistemology, see McPherson and Plunkett (2021).
55 Kornblith (2002, 1).
56 Which questions will be on the list of topics here (e.g., ones about nature, ground, supervenience, etc.) obviously interacts with broader issues in metaphysics and meta-philosophy about which of these topics are legitimate and important ones to investigate. We don't take a stand on these issues here.
57 See McPherson and Plunkett (2021), drawing on McPherson and Plunkett (2017) and Plunkett and Shapiro (2017).

58 We take our sympathy to the value of a pluralism of approaches to inquiry (in epistemology and elsewhere), as sketched here, to be more extensive than Kornblith's concession that there is room for a kind of "division of labor" in epistemic inquiry (Kornblith 2013, 213).

59 For further discussion of different ways in which this kind of "meta"-level project in epistemology might be instrumentally useful to the kind of "object-level" inquiry we are considering here (and other ways in which the projects can connect to each other), see McPherson and Plunkett (2021; 2017); Plunkett and Shapiro (2017); McPherson and Plunkett (2025), and McPherson (2012).

60 See McPherson and Plunkett (2021), drawing on McPherson and Plunkett (2024a) and McPherson (2012).

References

Ballantyne, N. (2019). *Knowing Our Limits*. Oxford: Oxford University Press.

Bishop, M.A. and Trout, J.D. (2004). *Epistemology and the Psychology of Human Judgment*. New York, NY: Oxford University Press.

Boyd, R. (1988/1997). How to be a moral realist. In: *Moral Discourse and Practice* (ed. S. Darwall, A. Gibbard, and P. Railton), 105–136. New York, NY: Oxford University Press.

Burgess, A. and Plunkett, D. (2013a). Conceptual ethics I. *Philosophy Compass* **8** (12): 1091–1101.

Burgess, A. and Plunkett, D. (2013b). Conceptual ethics II. *Philosophy Compass* **8** (12): 1102–1110.

Burgess, A. and Plunkett, D. (2020). On the relation between conceptual engineering and conceptual ethics. *Ratio* **33**: 281–294.

Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

Cappelen, H. and Plunkett, D. (2020). A guided tour of conceptual engineering and conceptual ethics. In: *Conceptual Engineering and Conceptual Ethics* (ed. A. Burgess, H. Cappelen, and D. Plunkett), 1–26. Oxford: Oxford University Press.

Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, vol. **42**. Oxford: Oxford University Press.

Eklund, M. (2017). *Choosing Normative Concepts*. Oxford: Oxford University Press.

Fantl, J. (2023). Guidance and mainstream epistemology. *Philosophical Studies* **180** (7): 2191–2210.

Fantl, J. and McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Gettier, E.L. (1963). Is justified true belief knowledge? *Analysis* **23**: 121–123.

Gilovich, T. (2001). *How we Know What isn't So: The Fallibility of Human Reason in Everyday Life*. New York, NY: The Free Press.

Haslanger, S. (1999). What knowledge is and what it ought to be: Feminist values and normative epistemology. *Philosophical Perspectives* (13): 459–480.

Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Noûs* **34** (1): 31–55.

Hetherington, S. (2018). *The Gettier Problem*. Cambridge: Cambridge University Press.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon.

Johnson King, Z. (2022). Radical internalism. *Philosophical Issues* **32** (1): 46–64.

Kornblith, H. (2002). *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.

Kornblith, H. (2012). *On Reflection*. Oxford: Oxford University Press.

Kornblith, H. (2013). Is there room for armchair theorizing in epistemology? In: *Philosophical Methodology: The Armchair or the Laboratory?* (ed. M.C. Haug), 195–216. Abingdon: Routledge.

Maguire, B. and Woods, J. (2020). The game of belief. *Philosophical Review* **129** (2): 211–249.

McPherson, T. (2012). Unifying moral methodology. *Pacific Philosophical Quarterly* **93** (4): 523–549.

McPherson, T. (2018). Authoritatively normative concepts. In: *Oxford Studies in Metaethics*, vol. **13** (ed. R. Shafer-Landau), 253–277. Oxford: Oxford University Press.

McPherson, T. and Plunkett, D. (2017). The nature and explanatory ambitions of metaethics. In: *The Routledge Handbook of Metaethics* (ed. T. McPherson and D. Plunkett), 1–25. New York, NY: Routledge.

McPherson, T. and Plunkett, D. (2020). Conceptual ethics and the methodology of normative inquiry. In: *Conceptual Engineering and Conceptual Ethics* (ed. A. Burgess, H. Cappelen, and D. Plunkett), 274–303. Oxford: Oxford University Press.

McPherson, T. and Plunkett, D. (2021). Conceptual ethics, metaepistemology, and normative epistemology. *Inquiry*.

McPherson, T. and Plunkett, D. (2024a). Metaethics and the conceptual ethics of normativity. *Inquiry* 67 (1): 93–126.

McPherson, T. and Plunkett, D. (2024b). Topic continuity in conceptual engineering and beyond. *Inquiry* 67 (9): 2847–2873.

McPherson, T. and Plunkett, D. (2025) After metaethics. *Philosophers' Imprint* 25 (22): 1–17.

McPherson, T. and Plunkett, D. (forthcoming). Authoritative normativity. In: Copp, D. and Rosati, C.S. (eds) *The Oxford Handbook of Meta-Ethics*.

McPherson, T. and Plunkett, D. (n.d.) Manuscript. Topic Continuity, Realism, and The Objects of Philosophical Inquiry.

Nolfi, K. (2020). Epistemically flawless false beliefs. *Synthese* **198** (12): 11291–11309.

Pappas, G. (2023). Internalist vs. Externalist conceptions of epistemic justification. In: *The Stanford Encyclopedia of Philosophy* [online] (ed. E.N. Zalta and U. Nodelman). Stanford, CA: Stanford University Press, spring 2023 edition https://plato.stanford.edu/archives/spr2023/entries/justep-intext.

Plunkett, D. and Shapiro, S. (2017). Law, morality, and everything else: General jurisprudence as a branch of metanormative inquiry. *Ethics* **128** (1): 37–68.

Pritchard, D. (2007). Recent work on epistemic value. *American Philosophical Quarterly* **44** (2): 85–110.

Risberg, O. (2023). Meta-skepticism. *Philosophy and Phenomenological Research* **106** (3): 541–565.

Rorty, R. (1980). *Philosophy and the Mirror of Nature*. Princeton, NJ: Princeton University Press.

Scharp, K. (2020). Philosophy as the Study of Defective Concepts. In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen and D. Plunkett, 396–416. Oxford: Oxford University Press.

Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.

Sosa, E. (2009). *Reflective Knowledge*. Oxford: Oxford University Press.

Thomasson, A.L. (2020). A pragmatic method for normative conceptual work. In: *Conceptual Engineering and Conceptual Ethics* (ed. A. Burgess, H. Cappelen, and D. Plunkett), 435–458. Oxford: Oxford University Press.

Weisberg, M. (2006). Water is not H_2O. In: *Philosophy of Chemistry* (ed. D. Baird, E. Scerri, and L. McIntyre), 337–345. Dordrecht: Springer.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review* **105** (4): 489–523.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.