

‘Knowledge’ as a natural kind term

Victor Kumar

Received: 1 September 2012 / Accepted: 9 April 2013 / Published online: 18 April 2013
© Springer Science+Business Media Dordrecht 2013

Abstract Naturalists who conceive of knowledge as a natural kind are led to treat ‘knowledge’ as a natural kind term. ‘Knowledge,’ then, must behave semantically in the ways that seem to support a direct reference theory for other natural kind terms. A direct reference theory for ‘knowledge,’ however, appears to leave open too many possibilities about the identity of knowledge. Intuitively, states of belief count as knowledge only if they meet epistemic criteria, not merely if they bear a causal/historical relation to the term. I will develop this objection and show that it is grounded in modal considerations central to Kripke’s work on reference. I will also argue that a more plausible externalist semantics for natural kind terms disarms the objection.

Keywords Knowledge · Natural kind term · Direct reference · Skepticism · Kornblith · Kripke

The view that knowledge is a natural kind promises a new approach in epistemology. Knowledge is to be studied empirically in broadly the same way that naturalists often study other psychological categories. This view about knowledge and its study has semantic implications that have not been vetted. The purpose of this essay is to defend the view that knowledge is a natural kind by articulating a plausible externalist semantics for ‘knowledge’ as a natural kind term.

According to the received externalist semantics owing to Kripke and Putnam, natural kind terms refer directly. A causal/historical relation, rather than a term’s descriptive meaning, determines what property the term refers to. Thus, if ‘knowledge’ is a natural kind term and refers directly, knowledge is just that which is uniquely common to those beliefs that bear a suitable causal/historical relation to the term.

V. Kumar (✉)

Philosophy Department, University of Arizona, 931 N Euclid Ave, #129, Tucson, AZ 85719, USA
e-mail: vckumar@email.arizona.edu

A direct reference theory for ‘knowledge’ is problematic because it leaves open too many possibilities about what knowledge is. For example, it seems to allow that states of knowledge could turn out to be false, so long as the relevant causal/historical relation picks out false beliefs. I will flesh out the objection, showing that it is grounded in modal arguments central to Kripke’s work on reference, but I will also argue that the objection can be disarmed.¹

In Sects. 1–3 I will characterize the methodology and semantics appropriate to the view that knowledge is a natural kind and make a *prima facie* case that ‘knowledge’ behaves in epistemological discourse as prototypical natural kind terms behave in scientific discourse. In Sects. 4–6 I will turn to the objection sketched above. The objection targets a direct reference theory for ‘knowledge,’ but a recent and independently more plausible externalist theory of reference for natural kind terms is safe from the objection. Semantic and methodological issues are closely linked in my treatment, and by defending a semantics for ‘knowledge’ I will also advance the methodology of naturalized epistemology.

1 Knowledge as a natural kind

Kornblith (2002) has proposed that epistemologists model philosophical investigation of knowledge on scientific investigation of natural kinds. Kornblith formulates his view in distinctive terms, but arguably other naturalists like Dretske (1982) and Millikan (1993) also think of knowledge as a natural kind. For these philosophers knowledge is not an abstract and *a priori* accessible universal, but instead a mental state with concrete psychological causes and effects, of which empirical study is possible. In this section I will rehearse the grounds for this naturalized view in epistemology and explicate the affiliated methodology. My focus in the rest of the essay is on its semantic commitments.

Kornblith contends that knowledge is a natural kind by appealing to its place in scientific study of the mind. He and Williamson both argue that knowledge participates in psychological explanations of thought and action and does so ineliminably; explanations that cite knowledge cannot instead cite belief, or true belief, or even justified true belief without loss in explanatory power (Kornblith 2002, Chap. 2; Williamson 2000, pp. 7–8, 60–63, 75–76). Kornblith draws a different conclusion from this evidence, but his work and Williamson’s are otherwise complementary. Synthesizing their work yields a more complete causal/explanatory profile for knowledge.

Kornblith argues that knowledge is a legitimate object of scientific investigation. Of central interest to Kornblith is that cognitive ethologists regularly study knowledge in animals (Kornblith 2002, pp. 52–55). Chimpanzees, dolphins, birds, and many other non-human animals are often characterized as knowing many things about their physical and social environments. The knowledge they possess explains how they successfully meet their needs and interests, and thus too how their ancestors survived

¹ Variations on the objection often arise in discussion, but it has not been carefully developed in the literature (cf. Weinberg 2006, pp. 59–60). I first heard a version of the objection in conversation with Stewart Cohen.

and reproduced (pp. 56–61). For example, take Ristau’s research on piping plovers (Ristau 1991, cited in Kornblith 2002, pp. 53–55). As Ristau tells it, the birds know that certain intruders in their territory are dangerous, they know the location of the intruders, and they know the location of their young. Plovers use this knowledge, Ristau says, to feign injury and lead potential predators away from their nests. Knowledge explains their successful goal-directed activity and thus too the birds’ evolutionary fitness.

Williamson argues that explanations of human action often pointedly appeal to knowledge. Someone found the hidden treasure because he knew its location (Williamson 2000, pp. 61–62). Not only did he merely attempt to dig it up, but he found the treasure because he knew where it was buried. In another case Williamson tells us that a burglar spent all night ransacking a house because he knew there was a diamond hidden somewhere there (p. 62). If he only believed truly, he might have given up sooner. But because he knew that a diamond was there he spent all night searching.

For both Kornblith and Williamson, knowledge is ineliminable from explanations of successful action. It is because the explanandum is not just action simpliciter, but successful action, that the explanatory role of knowledge is not exhausted by the explanatory role of belief. Granted, whether or not a belief is knowledge—whether the belief is true, justified, based on evidence, reliably produced, etc.—has no immediate effect upon a subject’s physical movements. But once the explananda are individuated more broadly so as to encompass a subject’s action, how she acts in light of the feedback she receives, and the effects of her action upon her interests, knowledge gains explanatory power over belief.

First, true belief is more likely than false belief to lead to successful action, other things being equal, because it enables one to capitalize on genuine causal relations that are instrumentally relevant to the satisfaction of desires and interests (cf. Kornblith 2002, pp. 56–58). Second, as Williamson argues, knowledge has explanatory power because action is typically an extended process. Knowledge is more likely than true belief to lead to successful action because one is more likely to retain knowledge in the face of countervailing evidence (Williamson 2000, pp. 7–8, 61). Thus, knowledge is explanatorily more powerful than belief because justification, like truth, increases the probability of successful action.²

Kornblith and Williamson agree that knowledge is a mental state. Because it is ineliminable from our best explanatory model of the mind, Kornblith argues, knowledge meets general criteria used to assess whether objects of study are natural kinds. Knowledge is a natural psychological kind, that is, because it explains thought and behavior. Thus, for Kornblith, a theory of knowledge is an account of a natural phenomenon, rather than, say, an internal concept of knowledge or a Platonic universal.

Once re-conceived as an empirical claim about the identity of a natural kind, a philosophical account of knowledge becomes a reduction of a kind of mental state to its

² Williamson might not put the point in terms of justification. But it is plausible that it is in virtue of something like justification—or, perhaps, whatever turns true belief into knowledge—that knowledge is more resilient than true belief in the face of countervailing evidence. See Williamson (2000, pp. 62–63) for why knowledge is explanatorily more powerful than (Gettierized) justified true belief.

constitutive psychological properties. Just as naturalistic philosophers study empirically the nature of belief (Fodor 1975), desire (Timothy 2004), intention (Holton 2009), emotion (Griffiths 1997), moral judgment (Kumar 2013a,b), and so on, naturalized epistemologists can study knowledge empirically. These projects in the philosophy of mind and elsewhere are obvious analogues, but so too is the naturalistic project to identify the good and the right with natural properties (Railton 1986a,b; Boyd 1988; Brink 1989). Knowledge is a mental state, but a reduction of knowledge must also account for its normativity (Kornblith 2002, Chap. 5; see also Railton 1989).

For those who follow Kripke, true identity statements are metaphysically necessary, whether they are analytic or synthetic. Thus, like an analysis of the concept of knowledge, a synthetic reduction of knowledge is necessary, if true. However, unlike an analysis, a synthetic reduction is knowable only a posteriori rather than a priori. The account must be tested against empirical investigation of knowledge itself.

Epistemologists must draw on empirical investigation of the psychological properties that are common to paradigm cases of knowledge and that differentiate them from other mental states—properties that enable knowledge to figure in explanations of the sort that Kornblith and Williamson describe. A successful reduction of knowledge will capture the actual cases of knowledge that invite epistemological theorizing, but it will also yield greater understanding of why knowledge explains successful action.

In sum, the naturalized approach advanced by Kornblith pursues a traditional philosophical question: what is knowledge? If knowledge is a natural kind, however, an answer to this question turns on empirical criteria of evaluation of the sort that apply to synthetic reductions of other natural kinds. My interest in this essay is in the consequent semantics for ‘knowledge,’ and it is that to which I now turn.

2 How ‘knowledge’ refers

The naturalized methodology suited to the view that knowledge is a natural kind is reflected in a corresponding semantics. If knowledge is a natural kind, then the satisfaction conditions for ‘knowledge’ cannot be discovered through armchair reflection of the sort that is characteristic of traditional conceptual analysis. Nor is experimental investigation of the folk concept of knowledge adequate. What ‘knowledge’ refers to is fixed at least in part by the world—by actual cases of knowledge—and not by the content encoded in an internally represented concept of knowledge. In this section I will outline an externalist semantics for ‘knowledge’ as a natural kind term, distinguishing it from an internalist semantics that plausibly underwrites conceptual analysis.

Arguably, epistemologists who practice conceptual analysis assume a descriptivist account of the meaning of epistemic terms, if not explicitly then implicitly, as reflected in their methodology. According to descriptivism, ‘knowledge’ is defined by a uniquely specifying description: that one has a true belief, perhaps that the belief is based on evidence, perhaps also that the strength of one’s evidence co-varies with the practical significance of the proposition believed (either to the attributer or the subject), and so on. Someone knows a proposition if and only if she satisfies this complex definition. An implicit grasp of the definition underlies our ability to intuit whether individuals

have knowledge in hypothetical cases. The goal of epistemology, so conceived, is to reveal the meaning of ‘knowledge’ implicit in use. The meaning of ‘knowledge’ is its intension and *intension* determines *extension*.³

According to the externalist semantics developed by Kripke (1980), Putnam (1975), Devitt (1981), Scott (2002) and others, natural kind terms refer directly, bypassing intension and thus without being mediated by an internally represented definition. If ‘knowledge’ is a natural kind term, its reference is fixed initially, and perhaps repeatedly, by ostending to, or being in causal contact with, or being causally regulated by paradigm instances of knowledge, where experts’ judgments may determine what counts as a paradigm instance (see Kornblith 2002, pp. 11–12). ‘Knowledge’ then refers to those paradigm instances along with whatever else is a member of the same natural kind. Reference is subsequently passed on to other speakers via a causal chain, or perhaps a chain of communicative intentions.

Variations on a causal/historical account of reference are possible. What is essential to any externalist view is the negative thesis that a natural kind term does not refer mediately via a definition; reference of word to object is external to “meaning” in the sense of common, internal, descriptive definition. Later on in the essay, I will expound an externalist semantics that is an alternative to a pure direct reference theory. According to this view, a natural kind term refers to the property (1) that bears a causal/historical relation to the term and (2) that supports the causal/explanatory role of the kind in empirical generalizations. For the purposes of this essay I will assume that one or another externalist semantics for natural kind terms is correct. My task is to defend an externalist semantics for ‘knowledge,’ not to defend an externalist semantics for natural kind terms in general.⁴

In order to evaluate an externalist semantics for ‘knowledge,’ I will focus primarily not on knowledge itself, but on our discourse about knowledge. For ‘knowledge’ to be a natural kind term, I will assume, is for it to behave like one in epistemological discourse, for philosophers to be disposed to use ‘knowledge’ in relevantly the same way that scientists use prototypical natural kind terms.⁵ In general—and this is revealed by well

³ In general, descriptivist accounts vary along a number of dimensions, most of which are not directly relevant to the present topic. For example, the accounts may be offered as theories of meaning, or only as theories of reference; the description may specify either necessary and sufficient conditions or a cluster of conditions, perhaps weighted, none of which is necessary; classic descriptivism also differs from causal and rigidified descriptivism (see fn. 4).

⁴ Some philosophers argue that natural kind terms refer not directly but via causal descriptions (Lewis 1984) or rigidified descriptions (Lewis 1984). Causal and rigidified descriptivism are externalist theories, according to my characterization of externalism, because causal and rigidified descriptions do not provide definitions for ‘knowledge.’ Neither causal or rigidified descriptivism about ‘knowledge’ can support an alternative conception of epistemology as conceptual analysis. Epistemologists trading intuitions about lotteries, bank cases, fake barns, clairvoyants, etc. are plainly not trying to uncover a causal or rigidified description. The arguments I provide for an externalist account of ‘knowledge’ (Sect. 3) and the objection I develop (Sect. 4) apply as much to these views.

⁵ An anonymous reviewer asks why I restrict attention to philosophers’ usage. There are really two questions here. First, why restrict attention to usage in a technical discipline, rather than ordinary usage? My answer is that I do not necessarily think that ‘knowledge’ does, or should, function like a natural kind term in ordinary discourse. The “folk” may well treat ‘knowledge’ and other natural kind terms as non-kind terms. Furthermore, the aims of ordinary discourse are different from those of theoretical discourse, and it may well be that the aims of ordinary discourse are not served by treating ‘knowledge’ as a natural

known thought experiments to be discussed presently—this entails that speakers are disposed to let the extension of a natural kind term be fixed by something other than its definition. In the next stage of the essay I will suggest that ‘knowledge’ behaves like other natural kind terms—without yet arguing that one or another externalist view of reference best explains this behavior. So, the thesis that ‘knowledge’ is a natural kind term is for now sufficiently general that we do not yet need to take a stand on certain finer points of externalist semantics that will turn out to be of importance later on.

Finally, notice that the thesis that ‘knowledge’ is a natural kind term does not entail that knowledge is a natural kind. In general, an expression can behave in discourse like a natural kind term even if it does not pick out a natural kind. For example, ‘phlogiston’ and ‘jade’ once behaved like natural kind terms—in the sense that speakers were disposed to let their extensions be fixed by something other than their definitions—even though it turns out that phlogiston and jade are not natural kinds. So, it is possible that ‘knowledge’ is a natural kind term even if it turns out that knowledge is not a natural kind, either because there is nothing there for ‘knowledge’ to refer to (like ‘phlogiston’) or because the term denotes a class of disparate objects (like ‘jade’).⁶

3 A prima facie case

In this section I will make a prima facie case for the view that ‘knowledge’ is a natural kind term. My aim is to show that this view has certain advantages over the descriptivist semantics for ‘knowledge’ characterized above.⁷ I will begin by rehearsing criteria from Kripke and Putnam that distinguish natural kind terms from linguistic expressions that behave in the way that descriptivism predicts. I will then argue that these criteria, when applied to ‘knowledge,’ suggest that it too behaves like a natural kind term.

An externalist semantics for natural kind terms in science has several well-known advantages over a competing descriptivist theory, only two of which are yet important for my purposes. First of all, externalism explains how reference is secured even when individuals are mistaken, even grossly mistaken, about the nature of the referent. As Kripke argues, proper names can refer to their objects even when speakers associate only false descriptions with the names. Someone refers successfully to Einstein even

Footnote 5 continued

kind term. Second, given that I restrict attention to usage of ‘knowledge’ in a technical discipline, why focus on philosophers’ usage rather than scientists’ usage? My answer to this question is that what we should want to know is whether it is consistent with acceptable epistemological practice for philosophers to treat ‘knowledge’ as a natural kind term. Hence my focus in the second half of the essay on whether such treatment leaves open too many possibilities about the identity of knowledge. This should be of concern to epistemologists, but, I suspect, not of much concern to scientists.

⁶ Some philosophers hold that an expression is a natural kind term only if it refers to a natural kind (e.g., Wikforss 2005). However, the most powerful arguments for a direct reference theory depend on how natural kind terms are used, and our modal intuitions concerning them (see Sects. 3 and 4). These arguments depend on how we use language, not what our language refers to.

⁷ There may be yet other semantic theories for ‘knowledge,’ endorsed by those who reject Kornblith-style naturalized epistemology and traditional conceptual analysis, but to consider these theories would take us too far afield. Anyway, the arguments I will be adapting, from Kripke and others, themselves address only descriptivist theories as alternatives.

though the only uniquely specifying description she associates with ‘Einstein’ is that he invented the atomic bomb (Kripke 1980, pp. 84–85). Similarly, at one time those who used the term ‘water’ were quite mistaken about what is essential to water, thinking of it as a basic element. Even though H_2O fails to satisfy the descriptions associated with the term, ‘water’ in their dialect still refers to H_2O . Because ‘water’ bears a causal/historical relation to H_2O , an externalist theory, but not a descriptivist theory, can account for this kind of referential success.⁸

A second and related advantage concerns the ability of externalism to account for stability of subject matter across scientific revolutions (Putnam 1975, pp. 235–238; Boyd 1983). In scientific revolutions, the story goes, researchers working in a new paradigm use the same terms contained in the old paradigm but possess radically different definitions for the term. If descriptivism were correct, researchers in the old paradigm and researchers in the new paradigm would not disagree. Their words would have different meanings and, consequently, they would not be theorizing about the same thing. On an externalist theory, however, the old and new theories share the same subject matter because the natural kind terms in both theories bear a causal/historical relation to the same property.

These arguments should be familiar. What I want to argue is that ‘knowledge’ behaves like a natural kind term—behaves, that is, in just the ways that support an externalist semantics for natural kind terms in science over a competing descriptivist theory. First, individuals refer successfully to knowledge even when they are grossly mistaken about the nature of the referent. For example, some Platonists have thought mistakenly that to know is to recollect. Yet, intuitively, ‘knowledge’ in their dialect refers to knowledge. This is exactly what we should expect if ‘knowledge’ behaves for us like a natural kind term. We are ready to interpret Platonists’ uses of ‘knowledge’ as co-referential with ours even though the descriptions they associate with the term are grossly mistaken.

Second, that ‘knowledge’ is a natural kind term helps account for stability of subject matter across “epistemological revolutions.” In the fallibilist revolution epistemologists realized that knowledge does not entail infallible grounds for belief. More controversially, in the externalist revolution epistemologists realized that justification or warrant depends on features of beliefs, their interrelations, or their etiology that are inaccessible to consciousness. Prior to each revolution, the term ‘knowledge’ was defined by descriptions that genuine cases of knowledge fail to satisfy. If descriptivism were correct we would be forced to conclude that the fallibilist revolution was, rather than an advance in the theory of knowledge, a change in subject matter. For whatever reason, ‘knowledge’ came to mean something else that does not entail infallibility. But this is incorrect. Infallibilism was rejected, epistemological criteria for knowledge thereby made more permissive, because with the growth of science we managed

⁸ Kripke does not rehash his argument about proper names when he later discusses natural kind terms, as I have done here. Kripke’s central argument against descriptivism about natural kind terms in *Naming and Necessity* appeals to certain modal intuitions: for any description ‘F’ associated with the natural kind term ‘x,’ it might have been false that Fx. Intuitions like this are supposed to show that proper names and natural kind terms refer directly and not mediately via a description. In Sect. 4 I will develop an objection to the thesis that ‘knowledge’ is a natural kind term on the basis of the absence of analogous intuitions about ‘knowledge.’

to acquire considerable knowledge—beliefs similar to our paradigm instances—based on reasoning that is manifestly non-demonstrative.⁹

So far in this section, I have made a *prima facie* case for the view that ‘knowledge’ is a natural kind term. The view explains (1) how epistemologists refer successfully to knowledge despite having a grossly mistaken definition of ‘knowledge’ and (2) stability of subject matter across radical historical shifts in epistemology. In both cases, if we regard ‘knowledge’ as a natural kind term we can explain how radically divergent epistemological theories are linked. They share a common reference. In other contexts this kind of evidence is typically thought to support a direct reference theory, but I will eventually distill from the literature an alternative externalist semantics that is equally consistent with this evidence and that better accounts for the behavior of natural kind terms.

The view that ‘knowledge’ is a natural kind term is plausible, I have argued, but before I examine an apparent weakness in the next section I want to consider how descriptivists might respond. Descriptivists about ‘knowledge’ may object that the arguments above do not rule out the possibility that ‘knowledge’ refers mediately via a description. The objection turns on distinguishing the meaning of a term from a theory about its meaning, and runs as follows. Among the many descriptions associated with ‘knowledge’ is a core description that defines the term. Epistemologists who have a grossly mistaken theory about the nature of knowledge still implicitly possess the definition, even if they do not know it and judge incorrectly that ‘to know’ just means, say, to recollect. Furthermore, so-called revolutions in epistemology involve radical changes in the theory about what ‘knowledge’ means, but are not so radical that they involve changes to the definition. Core definitions, if not much else, are preserved across epistemological revolutions. Thus, according to this objection, descriptivism about ‘knowledge’ remains a tenable position, once a core description that defines the term is distinguished from descriptions that are part of false theories about the meaning of ‘knowledge.’

In this suggestion we find reason not to reject a descriptivist account of ‘knowledge’ prematurely; the rest of the essay is a defence against this possibility. However, absent specification of the core description that defines ‘knowledge’—one that must be sufficiently precise to pick out knowledge uniquely—the suggestion is difficult to assess. If the objection is to be compelling, proponents must show that the existence of a core description that defines ‘knowledge’ is more likely than the existence of a core description associated with natural kind terms in science that fixes their extension. Otherwise the objection merely challenges a general externalist semantics for natural kind terms and makes no special case for descriptivism about ‘knowledge.’

Although the suggestion cannot be ruled out in advance of concrete philosophical analysis of ‘knowledge,’ available inductive evidence counts against it. The pervasive-ness and intractability of disagreement among epistemologists suggests that there is no single definition of ‘knowledge’ that answers to all our disparate intuitions. The failure of epistemologists to agree on a solution to the Gettier problem is a case in point. Epistemologists disagree even about whether the concept of knowledge is a nor-

⁹ See [Kornblith \(2002, p. 17\)](#) for a complementary discussion of the semantics for ‘knowledge’ and the evolution of epistemology, though it is not couched in terms of epistemological revolutions.

mative concept. Notwithstanding all of this basic disagreement, epistemologists still manage to be theorizing about the same thing. A fair explanation is that an externalist semantics for ‘knowledge’ is correct. Whether it is an otherwise viable explanation will be assessed in what follows.

Admittedly, perhaps, epistemologists sometimes talk past one another. For example, it may be that some epistemologists are analyzing “animal” or “unreflective” knowledge, while other epistemologists are analyzing “human” or “reflective” knowledge (Sosa 1997; Lehrer 2000; cf. Kornblith 2002, Chaps. 3–4). Or it may be that the sort of justification that is thought to be a component of knowledge fractures into several different epistemic desiderata (Alston 2005). Possibilities like this would be actual, on the naturalized approach defended here, if it turns out that paradigm cases of knowledge lack unity. One possibility is that one class of paradigm cases does some of the explanatory work of knowledge and other classes do the rest. This remains an open empirical possibility. If it were to obtain, then we should conclude that knowledge is not one natural kind, but several. Empirical investigation of knowledge would fracture into two separate empirically grounded investigations.

I have suggested that basic disagreement about knowledge reflects different definitions. If the disagreement here is truly great—greater than in other philosophical debates—it is likely because philosophers lack a shared definition. In some cases the different definitions may denote distinct objects of study. Still, my conjecture is that the amount of basic disagreement about knowledge outstrips genuine differences in subject matter.

4 The objection from skepticism

In Sects. 1 and 2 I developed the epistemological approach that conceives of knowledge as a natural kind and ‘knowledge’ as a natural kind term. In Sect. 3 I rehearsed the case that natural kind terms in science do not refer mediately via definitions and argued that a parallel case can be made for ‘knowledge.’ The most familiar externalist semantics for natural kind terms is a direct reference theory. However, in this section I will develop an objection to the view that ‘knowledge’ refers directly. The “objection from skepticism” attempts to vindicate the descriptivist suggestion just floated, that there is a core description that defines ‘knowledge,’ possessed even by those who are grossly mistaken about what knowledge is and preserved across epistemological revolutions.

‘Knowledge’ satisfies several related criteria that distinguish natural kind terms from linguistic expressions that are amenable to a descriptivist treatment. First, someone refers successfully to a natural kind even though her definition is grossly mistaken and does not in fact correctly describe the kind. Second, researchers before and after a theoretical revolution refer to the same natural kind even though their definitions are radically different and have distinct extensions. A third criterion distinguishes natural kind terms, in what is arguably Kripke’s main argument against descriptivism in *Naming and Necessity*. This criterion, as I will now explain, suggests that ‘knowledge’ does not behave like a natural kind term.

Kripke argues that descriptivism is untenable because it cannot make sense of certain modal intuitions (1980, pp. 116–134). A descriptivist semantics for natural

kind terms would incorrectly rule out what seem to be genuine empirical possibilities, in which we discover that a natural kind does not fit our definition of the corresponding term. In one of his main examples Kripke asks us to suppose that ‘gold’ refers to whatever satisfies its definition, say, a yellow metal (pp. 116–119). It could turn out that none of our samples of gold have the properties in this definition. Observations that heretofore supported the definition have been mistaken. In that case, intuitively, gold is not what we thought it was. But descriptivism entails instead that gold does not exist, since nothing satisfies the definition. In general, Kripke argues, for any description by which competent speakers define a natural kind term, it could turn out that the corresponding natural kind does not satisfy the description. Therefore, descriptivism does not provide the right semantics for natural kind terms.

The objection from skepticism claims that this modal truth about natural kind terms is not true of ‘knowledge.’ The possibility of radical skeptical scenarios shows that it could *not* turn out that knowledge does not satisfy the definition of ‘knowledge.’ Specifically, it could not turn out that knowledge does not satisfy a definition that expresses overwhelmingly plausible epistemic criteria. I will now develop the objection in detail.

If ‘knowledge’ refers via a descriptive definition, then perhaps the most widely accepted definition is that knowledge is justified true belief (plus an anti-Gettier condition), where “justification” can be construed consistently with either internalism or externalism. (Nothing turns on this choice of definition, as we’ll see.) But now suppose that nearly all our beliefs are false, perhaps that we are brains in vats. In this skeptical scenario we do not have any knowledge—or at least not as much as we once thought. The problem with the view that ‘knowledge’ is a natural kind term and refers directly, it seems, is that it cannot make sense of this counterfactual about our lack of knowledge. On that view knowledge simply is whatever properties all or most of our paradigm cases of knowledge uniquely have in common—even if being false is among these properties. Thus, a direct reference theory for ‘knowledge’ must grant the possibility that genuine states of knowledge could be false beliefs. Descriptivism, by contrast, yields the desired outcome: if paradigm instances of knowledge fail to satisfy a core description that defines the term—that knowledge entails truth—then they are not genuine instances.¹⁰

Knowledge also entails justification on our best definition, and thus one can also run the objection from skepticism by imagining that our beliefs are unjustified. So long as all our paradigm cases of knowledge, besides being unjustified, share something else uniquely, a direct reference theory entails that knowledge is unjustified belief. But, it seems, knowledge simply could not be unjustified belief, and descriptivism respects this intuition.

Notice that one can run this objection using any favored definition of ‘knowledge’: imagine that paradigm cases of knowledge are unreasonable, or not based on evidence, or produced by unreliable processes, etc. In each case, it seems, natural kind theorists

¹⁰ Note the following technical difficulty. If all of our beliefs turned out to be false, then all or most cases of knowledge would have in common false warranted belief, but they would not share it uniquely. Thus, in the relevant radical skeptical scenario, paradigm cases of knowledge must have some other property uniquely, besides falsity.

must allow that reference can nevertheless succeed. The general point, transcending any putative definition of ‘knowledge,’ is that knowledge could turn out—quite implausibly—not to meet *any* plausible epistemic criteria embodied in definitions of ‘knowledge.’ Clear examination of the objection will be facilitated by focusing mainly on the first skeptical scenario, in which paradigm cases of knowledge are false beliefs.

I will spell out the implications of the objection in a moment, but I first want to dispense with an initial response. Naturalists may argue as follows. “Yes, I accept that knowledge could not be false belief. However, the description by which competent speakers allegedly define ‘knowledge’ also happens to express a necessary, a posteriori truth about the identity of knowledge. Consider, for example, Kornblith’s view that knowledge is reliably produced true belief. Suppose, just for the sake of argument, that this is the correct account of knowledge, justified empirically, and that it is an identity claim, true in all possible worlds. In that case it is necessary that states of knowledge are states of true belief. So, indeed, one cannot coherently imagine anything to the contrary. If knowledge *is* reliably produced true belief, it follows that knowledge could not be false belief. Thus, a direct reference theory for ‘knowledge’ in fact accords with the intuition that drives the objection.”

To properly understand the objection from skepticism, and to properly understand too why the foregoing response is inadequate, we must distinguish two kinds of modality, metaphysical and epistemic (Kripke 1980, pp. 123–125). It is metaphysically impossible that a natural kind does not satisfy descriptions that are part of a true identity claim. Nevertheless, it is epistemically possible that a natural kind does not satisfy descriptions that are part of a true identity claim. Something is epistemically possible, at least on one plausible account, only if it is consistent with our current evidence, as it seems to us. For example, given that cats are mammals, it is metaphysically impossible that cats are aliens. Nevertheless, it remains epistemically possible that cats are aliens (p. 143). It is consistent with our current evidence, as it seems to us, that what we call ‘cats’ could turn out to have a much different internal composition than we thought. The evidence, even from dissections, might have been misleading.¹¹

So, the objection from skepticism comes to this. If ‘knowledge’ is a natural kind term, it does not follow that it is metaphysically possible that states of knowledge are false beliefs. However, it does follow that this is epistemically possible—we could discover that states of knowledge are false beliefs, since it is consistent with our evidence that paradigm cases of knowledge are false beliefs. But, intuitively, it is not epistemically possible that knowledge is false belief. And a descriptivist account of ‘knowledge’ explains why in the strongest terms: it is conceptually impossible that states of knowledge are false belief. Knowledge, the descriptivist says, is that which answers to our definition of ‘knowledge,’ and one element of our definition is truth.

To be clear, the objection from skepticism does not claim that a direct reference theory for ‘knowledge’ completely rules out skepticism. Some possibilities remain in

¹¹ See Kripke (1980, pp. 103–104). In order to have a plausible conception of epistemic possibility, it is crucial to think about evidence in an experiential, qualitative way—evidence “as it seems to us.” Otherwise, if evidence is understood in terms of propositions with externally-fixed content, then our evidence may rule out too much. Similarly, an alternative account of epistemic possibility in terms of what is consistent with what we know might also rule out too much. We know that cats are not aliens, in fact, and thus in this sense the claim is not epistemically possible.

which natural kind theorists would correctly deny that we have knowledge. Consider by analogy two examples of natural kind terms mentioned earlier: ‘phlogiston’ and ‘jade.’ First, phlogiston is not a natural kind because there is no such thing as phlogiston—there is nothing out there in the world for ‘phlogiston’ to refer to. Analogously, we might come to accept eliminative materialism of a sort that denies the existence of all intentional states, knowledge included. In that case, while we thought we had defined ‘knowledge’ by appeal to paradigm cases, the existence of those cases was illusory. Hence, knowledge would not exist and skepticism would be true.¹² Second, jade is not a natural kind because there is no single kind for ‘jade’ to refer to. Paradigm instances of jade in fact consist of two different kinds of minerals, jadeite and nephrite. Analogously, if we discovered that there is nothing deep that unifies paradigm instances of what we take to be knowledge, we might conclude that there is no such thing as knowledge. There are instead a group of unrelated things, a gerrymandered kind, that we only call “knowledge.” Here too skepticism might follow.¹³

Natural kind theorists can therefore admit the truth of skepticism, if there is either nothing or no unified class of mental states to which the term ‘knowledge’ can refer. The objection from skepticism exploits a different kind of skeptical scenario. There does exist something for ‘knowledge’ to refer to, and it does make up a unified class. However, the members of the class are unified in part because they share a property that violates plausible epistemic criteria—they are all false (or unjustified, or unreasonable, or not based on evidence, or based on unreliable processes, etc.). The problem for natural kind theorists is that if ‘knowledge’ refers directly, then so long as all or most of our paradigm instances of knowledge are not illusory and have some underlying properties uniquely in common—even including being false—we cannot fail to have knowledge. But this is not epistemically possible. Thus, it seems, a natural kind approach to knowledge is mistaken.

5 An alternative externalist semantics

So far, the evidence I have canvassed is mixed. In some respects ‘knowledge’ seems to behave like other natural kind terms (Sect. 3), while in other respects it seems to refer mediately via a definition (Sect. 4). In this section I will explicate a recent externalist theory of reference for natural kind terms that accounts for scientific discourse better than a pure direct reference theory. In the next section I will use the theory to defend the thesis that ‘knowledge’ is a natural kind term against the objection from skepticism. But before that I will lay out my response to the objection in outline.

The objection from skepticism is an application of Kripke’s test designed to show that the semantics for natural kind terms leaves a great number of possibilities about

¹² The example is contentious. Eliminative materialism does not straightforwardly entail that the existence of paradigm cases of knowledge is illusory. For it is consistent with eliminative materialism that knowledge and its paradigm cases turn out to be brain states rather than intentional states. I will not try to develop the example so that it avoids this objection; its purpose is illustrative.

¹³ See Kornblith (2002, p. 23) for discussion. Certain Wittgensteinians will be drawn to a theory of knowledge that fits this characterization, on which instances of knowledge share not an essence and only a family resemblance, and hence might offer a non-skeptical interpretation.

the kind epistemically open. I will argue that ‘knowledge’ seems to fail the test only because of widely accepted but mistaken assumptions that Kripke and others make about the role that empirical theories play generally in fixing reference. My strategy will be to assimilate the problem cases for naturalism to scientific cases. In both cases paradigm instances of an alleged kind do share uniquely some underlying properties but, nevertheless, they are not in fact instances of the natural kind under investigation. The reason, in general, is that the underlying properties do not support the causal/explanatory generalizations in which the purported kind was thought, incorrectly, to participate. The kind does not exist, in short, because it cannot do the nomological work required of it. As I will explain, direct causal/historical relations between word and world are just one determinant of reference. The other determinant is causal/explanatory generalizations about the natural kind. When paradigm cases of an alleged kind do not support these generalizations—as in the skeptical scenarios described in Sect. 4, so I will argue—reference fails.

The semantics for natural kind terms that I will describe in this section is “mixed,” though still externalist. Theories of reference are mixed when they accord reference-fixing roles to causal/historical relations and descriptions. Many philosophers of language favor mixed theories of meaning or reference, e.g., two dimensionalist semantics (Jackson 1998; Chalmers 2006). (This view is mixed, but internalist.) Boyd (1999), among others (Griffiths 1997; Brigandt 2010, 2011), advances a mixed semantics for natural kind terms that is not simply a combination of direct reference theory and descriptivism. The view is externalist because it denies that the descriptions that help to fix reference are definitions possessed by all competent speakers. The descriptions are causal/explanatory generalizations and may be grasped only by experts.¹⁴

Boyd describes this view of reference in terms of “accommodation.” A natural kind must accommodate the inductive and explanatory practices associated with the corresponding natural kind term in order for reference to succeed. Boyd argues that this view explains why natural kind terms pick out genuinely projectable properties—in Goodman’s (1965) sense—rather than any number of other heretofore co-extensive “grue-like” properties. Rather than examine Boyd’s arguments in detail, I will offer further support for the view by appealing to evidence that is especially significant in the context of this essay, namely, cases of reference failure.

Some radical changes in theory preserve reference—in scientific revolutions, as I have already noted—but other radical changes in theory support abandonment of theoretical posits. Imagine, for example, the earnest empirical study of witches. Suppose that all paradigm instances of witches do share something uniquely in common. The old theory had it that witches are essentially the spawn of Satan, are endowed with magical powers and are malevolent. Imagine, however, that paradigm instances of witches are instead members of a natural biological sub-species of *Homo sapiens* with a distinctive genetic profile that causes red hair, anti-social personality disorder, and a condition that results in disfigurement and warts. The conclusion that investigators must draw is not that witches do exist and that the latter properties, or their

¹⁴ Of the authors cited, Brigandt has probably developed this sort of view in the greatest detail. However, Brigandt’s view is roughly that epistemic descriptions (for me, causal/explanatory generalizations), rather than helping to fix reference (as I believe), are an independent semantic dimension of a concept.

genetic bases, are essential to witches, but that there are no witches. The reason is not that the reference-determining meaning of ‘witch’ is the description mentioned above (being the spawn of Satan, etc.), but that the properties common to so-called witches do not support generalizations where witch once served as explanans. That someone is a witch was thought to explain why people in the village get sick, why husbands cheat on their wives, why crops are spoiled, etc. But it cannot be in virtue of having red hair, an anti-social personality and a condition that causes disfigurement and warts that these events occurred. Nevertheless, ‘witch’ may behave like a natural kind term. If paradigm instances of witches were discovered instead to be a sub-species of *Homo sapiens* that uniquely have in common some other set of properties that are able to support their causal/explanatory roles—not being the spawn of Satan, having magical powers, etc., but some natural, biological properties—we might conclude that witches do exist, though they are not what we thought they were.

The familiar externalist theory of reference for natural kind terms inherited from Kripke and others is as follows. A natural kind term refers to whatever empirical investigation reveals is uniquely common to all or nearly all those objects that are related directly to the term. However, this view is disconfirmed by cases in which we are forced to abandon theoretical posits. Some descriptions play an essential supplementary role in fixing reference—not descriptions that capture what we think is essential to the kind, but generalizations about the causal/explanatory relations in which the kind participates. A natural kind term refers to the property (1) that is uniquely common to some collection of objects related directly to the term and (2) that supports the kind’s causal/explanatory roles in empirical generalizations.¹⁵ Call this mixed externalist semantics for natural kind terms “MENK.” In the final section of the essay I will motivate this semantics for ‘knowledge’ by arguing that, unlike a direct reference theory, it is immune to the objection from skepticism.

Before I respond to the objection, however, I wish to linger on one important way in which MENK diverges from a direct reference theory.¹⁶ According to a direct reference theory, descriptions that speakers associate with a natural kind term play absolutely no role in fixing reference. Thus, it follows that speakers can be grossly mistaken about the nature of a kind but nevertheless refer successfully to it. According to MENK, descriptions about the kind’s causal/explanatory roles play a partial role in fixing reference. MENK thus entails there is one way in which speakers *cannot* be radically mistaken about the nature of a kind while referring successfully to it. We must be careful here. The description in question, about the explanatory roles of the kind, is not part of the definition of a natural kind term. Speakers can be competent with a natural kind term without possessing the relevant description. So, MENK still allows for radically mistaken definitions that are compatible with successful reference, and thus is compatible with the arguments rehearsed in Sect. 3 that are generally thought to support

¹⁵ A complete defence of this view would also rule out a causal or rigidified descriptivist analogue: a natural kind term refers to the property that satisfies a causal or rigidified description and that supports the kind’s causal/explanatory roles. I will not take up that defence here.

¹⁶ Thanks to an anonymous reviewer for pressing me on this point.

a direct reference theory. However, MENK entails that those who do possess the description—experts, perhaps—cannot be mistaken in their belief. For example, gold cannot but play the causal/explanatory roles assigned to it in chemical theory. Of course, it is possible that what we call “gold” does not play that role, but then our paradigm cases of gold would not in fact count as instances of gold. This is precisely what the arguments above in support of MENK have concluded.

6 Response to the objection

I turn now to application of the mixed externalist theory of reference described above. The objection from skepticism asks us to imagine that paradigm cases of knowledge turn out to be false or unjustified beliefs. This skeptical scenario is supposed to show that those who treat ‘knowledge’ as a natural kind term are committed to the epistemic possibility that knowledge does not entail truth or justification. However, I will now argue that in the imagined skeptical scenario, ‘knowledge’ would not refer to states of false or unjustified belief, given MENK, because these properties do not support psychological explanations in which knowledge enters.

As I made clear in Sect. 1, Kornblith and Williamson both argue that knowledge explains successful action. For example, knowing something about the size, shape and location of nearby objects explains successful navigation of one’s physical environment. By combining Kornblith’s and Williamson’s treatments, I argued that knowledge has explanatory power both in virtue of truth and justification. First, someone who has a true belief, rather than a false belief, is more likely to capitalize on instrumentally useful causal relations. Thus, I am more likely to arrive at a destination if I have a true belief about its location. Second, someone who has a justified true belief, rather than an unjustified true belief, is more likely to retain the belief and successfully carry out the action in the face of countervailing evidence. Thus, I am more likely to arrive at my destination if I know its location, even more so than if I merely have a true belief about where it is, because I am less likely to be steered in the wrong direction by misleading evidence.

It follows that if radical skeptical scenarios were to obtain, knowledge-based explanations of successful action would be debunked. If, as the critic imagines, paradigm cases of knowledge turn out to be false beliefs, “knowing” that P makes one no more likely to successfully carry out some action than does falsely believing that P. For one is no more likely than those with false beliefs to capitalize on the causal relations relevant to the satisfaction of intrinsic desires. Or, if paradigm cases of knowledge turned out to be unjustified beliefs, “knowing” that P makes one no more likely to achieve success than does merely truly believing that P. For one is no more likely to retain belief in the face of countervailing evidence.

According to MENK, a natural kind term refers to properties that are directly related to the term and that support the kind’s causal/explanatory role in empirical generalizations. Thus, if ‘knowledge’ is a natural kind term, it refers to properties that are directly related to ‘knowledge’ and that support the role of knowledge in explanations of successful action. If paradigm cases of knowledge are false belief,

then, as I have argued, they cannot support explanations of successful action.¹⁷ And thus, according to MENK, it would not follow that knowledge is false belief.¹⁸ That claim is not in fact epistemically possible.

In the course of unpacking the objection from skepticism in Sect. 4, I claimed that something is epistemically possible only if it is consistent with our current evidence, as the evidence seems to us. I have argued now that it is not epistemically possible that knowledge is false or unjustified belief. But in making this argument I am relying on a richer account of epistemic possibility. The idea is that what is epistemically possible is constrained not just by our current evidence, but also by the condition that the referent must support the causal/explanatory roles of the kind in empirical generalizations.

The arguments in this section and the last about reference do double duty as arguments about epistemic possibility. ‘Witch’ would not refer to a sub-species of *Homo sapiens* with red hair, anti-social personal disorder, etc., because the relevant properties cannot support explanations that cite witches. It could not turn out that witches, per se, are simply red haired, anti-social individuals with warts. Likewise, and for the same reasons, it is not epistemically possible that knowledge is false belief. False beliefs would not count as knowledge because they cannot support explanations of successful action. Thus, the argument ‘knowledge’ would not refer to false belief also shows that the identity is not epistemically possible.

In the skeptical scenarios that drive the objection from skepticism, ‘knowledge’ would not refer to states of false or unjustified belief because these states would not meet a condition of successful reference (and constraint on epistemic possibility): they would not support the causal/explanatory roles of knowledge in psychological generalizations. But consider the following objection. Imagine a skeptical scenario in which an evil demon, besides leading us to have false beliefs, ensures that whenever we act on a false belief we succeed. In this scenario, it might seem, states of false belief *would* support the relevant explanations. Thus, on the externalist view of reference I have endorsed, MENK, it seems to follow that the natural kind term ‘knowledge’ would refer to false beliefs.

I believe that the objection is mistaken. We thought that knowledge explained successful action by causing behavior that capitalizes on instrumentally valuable causal relations and by making belief resilient in the face of misleading evidence. In this skeptical scenario states of so-called knowledge do not play this explanatory role. Their

¹⁷ To be clear, the relevant description that helps to fix reference, and that false belief does not satisfy, is that knowledge explains successful action, not that knowledge is justified true belief. It just happens that, as an empirical matter, justified true belief is what supports explanations of successful action. To illustrate, suppose we abandon the notions of truth and falsity, adopting instead, say, some graded notions of accuracy in belief. In that case, perhaps, something other than truth—though not falsity—will support knowledge-based explanations of successful action.

¹⁸ An anonymous referee offers the following alternative treatment, borrowing from Korman (2006). He/she suggests a view according to which what determines the reference of a term depends on which of a number of presuppositions obtain. Suppose, then, that it is a presupposition that ‘knowledge’ refers to paradigm cases of knowledge if those case are true and justified belief. Reference fails because this presupposition is false. My view is formulated in terms of reference-fixing descriptions rather than presuppositions, but I am not sure the difference is a deep one. However, any alternative to my view must not employ presuppositions that eliminate the possibility that people can be deeply mistaken about the nature of a referent but refer successfully to it.

influence on successful action is mediated by an evil demon. There is no nomological connection between false belief and successful action. It is possible that the evil demon has set up the world such that whenever someone has a false belief his or her action is successful. But this would not count as a genuine nomological connection, between false belief and successful action, because it would not be in virtue of the constitutive properties of the one that the other occurs.¹⁹ Put informally, it is really the evil demon that explains successful action; the false belief itself is not the operative factor. Thus, paradigm cases of knowledge do not in fact play the causal/explanatory role they once appeared to. Once again and for the same reason, reference fails.

If truth or justification is not essential to paradigm cases of knowledge, our definition of ‘knowledge’ will go the way of our definition of ‘witch’—but for empirical reasons of the sort that lead to abandonment of other theoretical posits, not because that is what our definition dictates. The view that ‘knowledge’ is a natural kind term thus does not entail that so long as all paradigm cases of knowledge have something in common, they are genuine instances of knowledge. If what is essential to so-called cases of knowledge cannot support their causal/explanatory role in empirical theories, we do not have knowledge at all.

To concede that descriptions partially determine the reference of ‘knowledge’ is not to back off from the naturalized approach that studies knowledge as a natural kind. For the descriptions in question are not those which a priori reflection on hypothetical cases is suited to uncover. The descriptions are instead empirical generalizations. A natural kind approach to studying knowledge must investigate not just the character of our paradigm instances but also causal/explanatory generalizations concerning knowledge. This is the most faithful way of modeling philosophical investigation of knowledge on scientific investigation of natural kinds.

Much work remains to be done in carrying out the positive research program that views knowledge as a natural kind. This essay has focused on the semantic commitments of this program, but I now want to offer some suggestions about how fellow advocates might pursue the first order project, in light of my arguments here. First of all, I have argued that naturalists are not constrained by typical a priori claims about knowledge. For example, even if there was an a priori connection between knowledge and evidence, as a matter of definition, if ‘knowledge’ does not refer mediately via a definition, then this connection does not impose a constraint on naturalized epistemological theorizing. However, descriptions about the explanatory role of knowledge do constrain what knowledge is. The specific content of these descriptions is an empirical matter, but it may well be a priori knowable that natural kinds must support some explanatory role or another. If so, then there remains one a priori constraint, however loose, on epistemological practice.

If a priori resources are limited, then what empirical resources are available for naturalists to draw upon? In laying out the grounds for the view that knowledge is a natural kind in Sect. 1, I synthesized Kornblith’s and Williamson’s arguments that knowledge explains successful action. Williamson’s arguments are decidedly armchair. However,

¹⁹ I am assuming here a non-Humean account of scientific laws and nomological generalizations. They are not merely a matter of regularities but issue from powers inherent to the properties involved. Here is not the place to defend this non-Humean account, but it is grounded in [Armstrong \(1983\)](#).

it seems to me that they are not a priori. Having justified beliefs raises the likelihood of successful action because those with justified beliefs are less likely than those with unjustified beliefs to be steered in the wrong direction by misleading evidence. This insight issues from a sophisticated folk psychological reflection on human thought and action, and thus seems to be empirical, even though it is not scientific. In sum, then, study of knowledge as a natural kind can be informed not just by scientific research but also, perhaps, by armchair-empirical reflection. If armchair-empirical theorizing yields insights about knowledge and its link to action then it may earn a proper place in naturalized epistemology.

But scientific research is also relevant to the project. Kornblith himself looks to the study of animal knowledge and behavior in cognitive ethology. But this raises the question of whether human knowledge, and its explanatory role, might be different from animal knowledge. I suggested in Sect. 3 that this is a live possibility and consistent with the naturalized approach under discussion. It might be that human knowledge and animal knowledge are different natural kinds, perhaps two species within an overarching genus, each with a distinctive explanatory role. It would be useful, then, and fully consistent with Kornblith's work, if there were a scientific program of research on human knowledge that could inform epistemological investigation without, like cognitive ethology, threatening to change the topic. Naturalists should perhaps focus their efforts here.

7 Conclusion

I have argued in this essay that 'knowledge' behaves like a natural kind term and that it does not refer mediately via a description that is part of a common definition of 'knowledge.' Speakers refer successfully to knowledge even though their definitions are grossly mistaken and do not pick out the right extension. Researchers before and after an epistemological revolution both refer to knowledge, even though their definitions are not co-referential.

Furthermore, I have argued that 'knowledge' would not refer in the sorts of cases that force abandonment of theoretical posits. To treat 'knowledge' as a natural kind term is to admit that typical a priori assumptions about knowledge could be mistaken. But it does not entail that knowledge is whatever we take to be knowledge. Paradigm cases of knowledge, if they are to deserve the name, must support the causal/explanatory roles assigned to knowledge in our best empirical understanding of the mind.

The epistemological approach that conceives of knowledge as a natural kind is thus compatible with taking skepticism seriously. Knowledge is an important psychological category in part because it helps make sense of how we are able to successfully navigate the world around us. But if investigation of our mental states debunks these explanations, we will have reason to be skeptical about the existence of knowledge. If investigation vindicates these explanations, we will have reason to accept the existence of knowledge. This is how natural kind theorists can give skepticism its due.

Acknowledgments For helpful feedback on earlier drafts of this essay, I would like to thank Selim Berker, Richmond Campbell, Stewart Cohen, Juan Comasana, Ian Evans, Terry Horgan, Michael Hymers, Uriah Kriegel, Farid Masrouf, Shaun Nichols, Mark Timmons, Jonathan Weinberg, and two anonymous referees

for this journal. I would also like to thank audiences at The University of Arizona, Harvard University, and the Canadian Philosophical Association 2010 Congress. Work on the essay was supported by a fellowship from the Social Sciences and Humanities Research Council of Canada.

References

- Alston, W. (2005). *Beyond justification: Dimensions of epistemic evaluation*. Ithaca, NY: Cornell University Press.
- Armstrong, D. (1983). *What is a law of nature?*. Cambridge: Cambridge University Press.
- Boyd, R. (1983). On the current status of the issue of scientific realism. *Erkenntnis*, 19, 45–90.
- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 181–228). Ithaca, NY: Cornell University Press.
- Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New interdisciplinary essays*. Cambridge, MA: MIT Press.
- Brigandt, I. (2010). The epistemic goal of a concept: Accounting for the rationality of semantic change and variation. *Synthese*, 177, 19–40.
- Brigandt, I. (2011). Natural kinds and concepts: A pragmatist and methodologically naturalistic account. In J. Knowles & H. Rydenfelt (Eds.), *Pragmatism, science and naturalism* (pp. 171–196). Frankfurt am Main: Peter Lang Publishing.
- Brink, D. (1989). *Moral realism and the foundations of ethics*. New York: Cambridge University Press.
- Chalmers, D. (2006). Two-dimensional semantics. In E. Lepore & B. Smith (Eds.), *Oxford handbook of philosophy of language*. Oxford: Oxford University Press.
- Devitt, M. (1981). *Designation*. New York: Columbia University Press.
- Dretske, F. (1982). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Fodor, J. (1975). *The language of thought*. New York: Cromwell.
- Goodman, N. (1965). *Fact, fiction and forecast*. Indianapolis: Bobbs-Merrill.
- Griffiths, P. (1997). *What emotions really are: The problem of psychological categories*. Chicago: University of Chicago Press.
- Holton, R. (2009). *Willing wanting waiting*. Oxford: Oxford University Press.
- Jackson, F. (1998). *From metaphysics to ethics: A defence of conceptual analysis*. Oxford: Oxford University Press.
- Korman, D. (2006). What externalists should say about dry earth. *Journal of Philosophy*, 103(10), 503–520.
- Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.
- Kripke, S. (1980). *Naming and necessity*. Cambridge: Harvard University Press.
- Kroon, F. (1987). Causal descriptivism. *Australasian Journal of Philosophy*, 65(1), 1–17.
- Kumar, V. (2013a). *Empirical reduction of moral judgment*. Under review.
- Kumar, V. (2013b). *Moral judgment as a natural kind*. Under review.
- Lehrer, K. (2000). Discursive knowledge. *Philosophy and Phenomenological Research*, 60(3), 637–653.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, 62(3), 221–236.
- Millikan, R. (1993). *White Queen psychology and other essays for Alice*. Cambridge, MA: MIT Press.
- Putnam, H. (1975). The meaning of 'meaning'. In K. Gunderson (Ed.), *Language, mind and knowledge* (pp. 131–193). Minneapolis: University of Minnesota Press.
- Railton, P. (1986a). Moral realism. *Philosophical Review*, 95, 163–207.
- Railton, P. (1986b). Facts and values. *Philosophical Topics*, 14, 5–31.
- Railton, P. (1989). Naturalism and prescriptivity. *Social Philosophy and Policy*, 95, 151–174.
- Ristau, C. (1991). Aspects of the cognitive ethology of an injury-feigning bird, the piping plover. In C. Ristau (Ed.), *Cognitive ethology: The minds of other animals: Essays in honor of Donald R. Griffin* Hillsdale, NJ: Lawrence Erlbaum.
- Timothy, S. (2004). *Three faces of desire*. Oxford: Oxford University Press.
- Scott, S. (2002). *Beyond rigidity*. Oxford: Oxford University Press.
- Sosa, E. (1997). Reflective knowledge in the best circles. *The Journal of Philosophy*, 94(8), 410–430.
- Weinberg, J. (2006). What's epistemology for? The case for neopragmatism in normative metaepistemology. In S. Hetherington (Ed.), *Epistemology futures*. Oxford: Oxford University Press.
- Wikforss, A. M. (2005). Naming natural kinds. *Synthese*, 145(1), 65–87.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.