

Natural kinds as categorical bottlenecks

L. R. Franklin-Hall

Published online: 24 May 2014
© Springer Science+Business Media Dordrecht 2014

Abstract Both realist and anti-realist accounts of natural kinds possess prima facie virtues: realists can straightforwardly make sense of the apparent objectivity of the natural kinds, and anti-realists, their knowability. This paper formulates a properly anti-realist account designed to capture both merits. In particular, it recommends understanding natural kinds as ‘categorical bottlenecks,’ those categories that not only best serve us, with our idiosyncratic aims and cognitive capacities, but also those of a wide range of alternative agents. By endorsing an ultimately subjective categorical principle, this view sidesteps epistemological difficulties facing realist views. Yet, it nevertheless identifies natural kinds that are fairly, though not completely, stance-independent or objective.

Keywords Natural kinds · Realism · Anti-realism · Homeostatic property clusters · Classification · Categorization

1 Introduction

Categorization—the grouping of individuals or processes into general classes or kinds—is essential for both our theoretical and practical grasps on the world. Absent categorization, each encounter, scientific or otherwise, would be a one-off, and life unavoidably “just one damned thing after another.” At least for creatures like us, this would make rational action, let alone thought itself, impossible.¹

¹ More modestly, the author of *Classification Made Simple* observes that without categorization “shopping would be very difficult and time consuming” (Hunter 2009, p. 3).

L. R. Franklin-Hall (✉)
Department of Philosophy, New York University, 5 Washington Place, Manhattan, NY 10003, USA
e-mail: lrf217@nyu.edu

While categorization appears nothing short of mandatory, less compulsory is the application of one particular scheme of categories (i.e., a classification) to a given domain of individuals. For instance, on the broadly popular pluralist view (Chakravartty 2007; Dupré 1995; Khalidi 1998; Kitcher 1984) there are many alternative and equally valid ways of joining individuals into groups or kinds, with different schemes apt to different projects.

Yet whether or not pluralism turns out to be correct, classification is not supposed to be ‘anything goes,’ with every conceivable grouping tracking a so-called *natural kind*. Instead, a “very basic intuition” (MacLeod and Reydon 2012, p. 91) shared by philosophers, scientists, and lay-folk alike is that there is a significant difference between the categories offered up by the mature sciences (e.g., *bosons*, *C. elegans*) and those not in scientific circulation—from the wildly pathological (e.g., *animals born on a Tuesday* (Kornblith 1995)), or the scientifically defunct (e.g., *consumption*, *hysteria*), to the merely boring (e.g., *white things* (Mill 1882)). While the former are understood as true or correct, the latter are dismissed as faulty or defective in some important sense.

But just what sense is that? Answers to this question can take either a realist or an anti-realist cast. On the realist approach, as traditionally construed, the universe possesses a mind-independent natural-kind structure, such that whether a category or classification “carves at the joints” has nothing to do with its place “in human languages, conceptual schemes, biology, or anything like that” (Sider 2012, p. 5).² Thus we do well or badly, classification-wise, to the extent that our partitions track the kinds embedded in nature itself, and the pathological categories are those that in no way—even but through a glass darkly—match the world’s own.

Though some do, anti-realists need not repudiate in full the picture just described; they may maintain the existence of natural kinds, suitably understood, along with the associated claim that the correct categories reflect them. Where they must part ways with the realist is in denying that the identity of the natural kinds is *fully mind-independent*, such that which groups form natural kinds, and which do not, is, from a metaphysical point-of-view, completely independent of us, the categorizers. Thus anti-realists frame accounts according to which the natural kinds depend—whether explicitly or implicitly—on our aims, concepts, or cognitive capacities. By the lights of some pragmatist theories, for instance, “[t]here’s no higher standard to which our concepts are to answer than the efficient satisfaction of the purposes of inquiry; those purposes are set, not by nature, but by us” (Kitcher 2008, p. 119). The pathological classifications—and the kinds corresponding to them—will then fall short because they are, quite literally, good-for-nothing.

Which of these two broad approaches is correct? This question is difficult, and the goal of this paper cannot be to answer it conclusively. Instead, my aspiration is to improve anti-realism’s prospects by developing the most plausible version of the anti-realist approach that I can. In particular, I will argue that natural kinds be identified with ‘categorical bottlenecks’: categories that well serve not only us, but

² As will become clear, I use *realism* throughout this paper to indicate a commitment to the mind-independence of facts about which kinds are natural, and not to indicate an ontological commitment to universals, a further claim that some but not all advocates of realism (in my sense) will endorse.

also a variety of agents possessing any of a range of aims and capacities, a range that—though ultimately relative to us—extends well beyond those we presently enjoy. Though not every universe will possess categorical bottlenecks, ours does partly in consequence of its causal architecture.³ More specifically, because many features are clustered together in the individuals we identify, and because individuals possessing these features play many different roles in systems in which they are embedded, quite diverse purposes flow, as liquid to the neck of a pouring bottle, to the same categories and classification—to the natural kinds.

After more precisely characterizing both realist and anti-realist positions in Sect. 2, the paper proceeds in Sects. 3 and 4 to review two challenges facing theories of natural kinds, the first concerning the progress we have made in identifying them and the second our access to facts about them. Section 5 then presents the Categorical Bottleneck Account, while Sects. 6 and 7 argue that it can successfully navigate these challenges—thus speaking in its favor over realist views. Section 8 goes on to sketch additional advantages of the Categorical Bottleneck approach over alternative anti-realist proposals.

2 Distinguishing natural kind realism and anti-realism

Though just what makes an account of natural kinds properly *realist* is hardly straightforward, it is customary—and I believe well-motivated—to understand realist views as those maintaining the full objectivity or mind-independence of the natural kinds. In particular, the defining claim of natural-kind realism is this: *there are at least some mind-independent facts or truths about which kinds are natural and which are not*. Though the facts or truths that the realist judges mind-independent are most commonly put in terms of a kind's 'naturalness' (or lack thereof), they may be formulated in alternative locutions, such as the 'reality,' 'non-reality,' or 'artificiality' of a kind or kinds. For example, that *human races are not real*, or that *protons form a natural kind*, both express such facts. Equivalently, Plato's famous carving metaphor may be used to do so, as when it is asserted that *the classifications of the DSM-5 do not carve nature at the joints*. For the sake of expository simplicity I will exclusively use the language of 'naturalness' here, but translations into alternative terminology are straightforward.

Though it sometimes pays to explore natural kind realism *in general*—that is, in the broad form just characterized—it is more customary to consider realism *with respect to a particular domain* of individuals or processes, such as organisms, atoms, or mental states. For instance, a realist about natural kinds in the atomic domain would hold that there are mind-independent facts about which groups of

³ By appealing to our universe's causal architecture in this way, the anti-realist view I will develop exclusively concerns facts about which kinds are natural, and does not extend to the causal order itself (the mind-independence of which I am here agnostic). Naturally, keeping these two domains separate requires a conception of causal order such that facts about it do not themselves depend on facts about the natural kinds whose mind-independence is in question.

atoms form natural kinds, and which do not. Such a realist would also endorse natural kind realism *in general*, as its demands are strictly weaker.

For facts about which kinds are natural to be mind-independent is for those facts to be independent of all aspects of us—including our evaluative attitudes, conceptual schemes, and cognitive capacities.⁴ I submit that a particular theory of natural kinds satisfies this condition (and is thus properly realist) just in case its *fundamental categorical principle*—its general claim of the form ‘the natural kinds are...’—is both *content independent* and *status independent*. These represent different pathways by which facts about the natural kinds might depend on us.

A theory of natural kinds is *content independent* when its analysans—i.e., what it identifies as natural kinds—does not involve us in any way. For instance, a theory that takes the natural kinds as primitive, perhaps understanding them as aspects of *the structure of the universe* (e.g. Sider 2012), almost certainly satisfies this demand, as might one identifying the natural kinds with essences (e.g. Ellis 2001). On the other hand, a view that identifies natural kinds with *groups the distinguishing of which helps us to achieve our ends* (e.g., Kitcher 2007) transparently fails the content-independence requirement, as its fundamental categorical principle explicitly refers to features of the categorizer. Alternatively, some views fail the content-independence requirement only when the commitments that inform the intended interpretation of the fundamental categorical principle are themselves considered. For instance, accounts that tie the naturalness of a kind to the role of that kind (or its connected category) in explanation and induction (e.g., Boyd 1991) will be anti-realist if explanatory and inductive norms are themselves understood as correct or incorrect only relative to us. They would also be anti-realist, via a somewhat different route, if the particular explanations and inductions that a kind (or its connected category) must serve were to be selected *by us*, rather than being privileged by the agenda of the mind-independent universe. After all, different categories are suitable to different explanatory and inductive targets, a circumstance that (as will be considered more below) can yield mind-dependent—and thus anti-realist—natural kinds.

Somewhat more subtly, a theory of natural kinds is *status independent* just in case its fundamental categorical principle is understood as correct or incorrect mind-independently, rather than it having an ultimately subjective grounding. That subjective grounding might involve some feature of our innate conceptual architecture or the output of a circuitous process of cultural change. No matter what, status independence is lacking when *some feature of us*—in the simplest case, our tendency, whether conscious or otherwise, to distinguish natural kinds from unnatural kinds using a particular principle—makes the fundamental categorical

⁴ How does this approach apply to mind-dependent individuals (or processes), such as human mental states (which are constitutively dependent on us) or particles only created in supercolliders (which are causally dependent on us)? Though it might appear that my definition of realism would—rather unfortunately—make natural kinds of such mind-dependent individuals impossible, this is a mirage. What must be mind-independent for the natural kind realist are facts concerning *which groups form, or do not form, natural kinds*, not the existence of the members of those groups.

principle true. In so doing, we make it the case (in concert with features of the mind-independent universe) that particular kinds are natural or not natural—a squarely anti-realist consequence.

Even granting the cogency of these two standards, it is often arguable whether a particular proposal satisfies them and is thus properly realist. One difficulty stems from the fact that a theory's status-independence cannot be read off of its fundamental categorical principle presented in isolation. For instance, take a broadly causal theory of natural kinds⁵—such as Richard Boyd's influential suggestion that natural kinds in some domains are groups of individuals that possess clusters of mutually stabilizing or sustaining properties (i.e., 'homeostatic property clusters').⁶ Such a theory may be a status-independent theory, but it could equally well be status-dependent. For this to be the case, the theory's fundamental categorical principle would need to be understood as correct in virtue of something about us—our concepts, aims, or practices. Another challenge comes in distinguishing a proposal's fundamental categorical principle—its assertion that *the natural kinds are such-and-such*—from more modest, but still informative, remarks about natural kinds that fall short of identity claims. To illustrate again with the help of Boyd's proposal: rather than maintaining that natural kinds *are* homeostatic property clusters (HPCs), it might be suggested that natural kinds *are often* HPCs, while being strictly identical with something else—optimally some feature that would help explain just why natural kinds were sometimes HPCs. For instance, Boyd seems to endorse such a strategy when he suggests that natural kinds are groups corresponding to categories that support a successful 'accommodation' between "our classificatory and inductive and explanatory practices" and the universe's causal structures (Boyd 1999a, p. 69; see also Boyd 1991). This proposal, if taken as a candidate fundamental categorical principle, can be used to explain why natural kinds *are often* HPCs, while also helpfully allowing that some natural kinds (perhaps those of physics or of chemistry) will not be. At the same time, the view is anti-realist in virtue of making the natural kinds depend on *our scientific projects and practices*. On such a theory the natural kinds would have been different had we developed different projects and practices best accommodated by different categories.

Even putting these complex interpretive issues aside, the approach to realism just described leaves a number of questions open, questions on which particular realist theories may differently commit. For instance, though realists tend to be categorical monists, the realism considered here accommodates both monist and pluralist views. (Very briefly, these differ on whether there are thought to be multiple, possibly cross-cutting, types of natural kinds of a particular sort, such as multiple varieties of biological species—e.g., genealogical and morphological.) Another choice-point for

⁵ By 'causal theory' I mean to pick out any theory that understands natural kinds as groups enjoying some particular kind of lower-level causal architecture rather than to refer to the 'causal theory of reference' articulated by (among others) Saul Kripke. (Needless to say, a causal theory of natural kinds is perfectly compatible with the correctness of the causal theory of reference for particular natural kinds, such as gold, water, swan, etc.)

⁶ See Boyd (1990, 1991, 1999a, b, 2000, 2010).

the realist concerns the thesis of *essentialism*, the view that members of a natural kind share a *real essence*—perhaps a locus of de re modality. As I understand things, essentialism is not a requirement of realism itself, but rather reflects a particular way of spelling out the realist vision.

With realism so defined, anti-realism is just the denial of realism, either *in general* (yielding a global anti-realism) or *with respect to a particular domain* (yielding a local anti-realism). And just as there are different ways of spelling out the realist position, so will there be different varieties of anti-realism. For instance, anti-realists like Ian Hacking (2007) and Nelson Goodman (1978)—who might also be labeled natural kind *nihilists*—completely reject the notion of a kind's naturalness, choosing to talk simply of 'relevant' or 'irrelevant' kinds when called upon to distinguish the interesting categories from the pathological overflow. Yet it is equally an option for the anti-realist to affirm the existence of natural kinds—maintaining that there *are* facts about which groups form natural kinds and which do not—but to take these facts to be partially mind-dependent. As an example, on one development of Kitcher's (2001, 2007) pragmatist view, the natural kinds correspond to those categories that would best serve the purposes, whether epistemic or practical, that our community would endorse at the end of a tutored democratic decision process.⁷ Anti-realists may also differ as to which aspects of us (in addition to the mind-independent universe) natural kinds depend on. After all, almost every anti-realist will affirm that facts about which kinds are natural are independent of what we *explicitly believe* the natural kinds to be. But whereas some pragmatist anti-realists will say that these facts may turn on our interests, aims or commitments—whether epistemic or practical—others will relativize the kinds to a restricted range of epistemic aims.

By understanding realism in terms of the mind-independence of facts about natural kinds, as I have suggested, realism and anti-realism offer radically contrasting pictures of the place of natural kinds in our universe. That contrast can be made vivid by considering what the natural kinds would have been had everything about our universe been the same as now—including the particular heavenly bodies, the major terrestrial geographical features, and the physical particles—except that humans had developed somewhat different characteristics. No matter the human difference envisioned, the realist's natural kinds—in all non-human domains anyway—would be the same as they actually are. But for the anti-realist the natural kinds *could* be different in both intension and extension, though the new form taken would depend on the extent of the human alteration envisioned. In the extreme, though, differences could extend as far as the natural kinds of physics. For a scenario in which such differences are grounded in the *content* (rather than the *status*) of an anti-realist's categorical principle, suppose it were true that, in a world where we had much greater cognitive powers, our actual classification of physical particles would make it easy for anyone to build horribly destructive weaponry. In that case, a pragmatist theory that relativized the natural kinds to both

⁷ There are really two strands in Kitcher's recent work on this topic, one maintaining that there are *no* natural kinds, and another suggesting that *there are* natural kinds, but that they are *relative to our ends*. I emphasize the second strand here because it presents a more interesting contrast to my own position.

our epistemic and practical aims might insist that the best classification—and thus the natural kinds—would be different in a way that would avert that nightmarish consequence (for a related scenario, see Kitcher 2007, p. 307).⁸

3 The anti-realist's problem: progress

Some will judge the feature of anti-realist views just highlighted—the subjectivity of their natural kinds—reason enough to opt for realism. After all, as will be explored more in Sect. 5, the idea that the (non-human) natural kinds are in any way reflections of us is rather counter-intuitive. Anti-realism also appears at a disadvantage when faced with a closely related phenomenon: the progress of science. That progress seems to consist not just in the development of increasingly accurate dynamic theories, but also in the elaboration of ever more correct classifications of the individuals that those theories describe. For instance, alongside Darwin's major achievement—the theory of evolution by natural selection—he offered a new view of the 'Natural System' of living things, suggesting that "propinquity of descent" was "the bond, hidden as it is by various degrees of modification, which is partially revealed to us by our classifications" (Darwin 1859, pp. 413–414).⁹ The adoption of Mendeleev's periodic table—a scheme often said to "reflect the natural order" (Scerri 2011, p. xix)—is an even better example of categorical progress due to the fact that it was proposed independently of a substantial advance in our understanding of the behavior of the materials categorized. Hence, its acceptance cannot so easily be seen as a downstream consequence of straightforward empirical progress.

⁸ In suggesting that, for the anti-realist, the natural kinds might have been different had we been different, I have made an assumption about how to interpret this counterfactual that is worth noting. I have assumed that for the anti-realist, the natural kinds are picked out, in a given world, by those groups that play the 'natural kind role' in that world, a role that anti-realists—whether promoting a content-dependent theory or one that is only status-dependent—will claim is partially determined by features of the categorizers in that world. Therefore, I am assuming that we are not availing ourselves of two rather squirrely ways of interpreting the counterfactual: (1) holding our present aims and capacities fixed by packing them into the 'natural kind role'; or (2) treating the term 'natural kind' as a rigid designator by tying its content to whatever physical property (if any) happens to play the natural kind role in the actual world. Having done either of these things, views I would consider to be anti-realist might maintain that, even had we been different, the natural kinds would have been the same. However, as David Lewis put it in a related discussion of laws of nature, this is "a cosmetic remedy only. It doesn't make the problem [of mind-dependence] go away, but only makes it harder to state" (1994, p. 479). For more in-depth treatments of this issue, played out within discussion of normativity and metaphysics, see Street (2006: §7; 2008; ms.) and Sider (2012: §4.4) respectively.

⁹ Advocates of the species-as-individuals thesis (e.g., Hull 1978; Ghiselin 1997), who will deny that biological taxa are natural kinds and will instead consider them to be natural individuals, can re-read all of my claims about species classification as those about individuation. This is possible because, though this paper targets realism about *natural kinds* for reasons of space, identical issues face realism about *natural individuals*; the realist claims that there are mind-independent facts about what the natural individuals are, while the anti-realist maintains that any such facts are relative to us. (See Boyd 1999a, p. 73 for congenial comments on this matter.)

Categorical progress seems to speak in favor of realism because the realist can so straightforwardly say what it consists in. In particular, for the realist, categorical progress involves the adoption of classifications that with increasing accuracy mirror nature-made groupings, the *natural kinds* (Psillos 1999, p. 280). Obviously, the anti-realist cannot endorse a mirroring analysis. Consider, for example, the discouraging situation of the anti-realist conventionalist: she maintains that the natural kinds depend on our “arbitrary choice” to endorse one scheme of categories over another (Mill 1882: bk 4, Ch. 3; Sider 2012, p. 54), such “that we could have proceeded differently” without going in any way wrong (Elder 2004, p. 8). On this view, categorical change may be equivalent to changing the side of the street on which traffic travels—not better or worse, but merely different. This picture would seem to render categorical progress an illusion, a threat that, according to some interpreters, drove even John Locke into realism’s embrace (Kornblith 1995, p. 18).

Of course, for this to be a triumph for realists assumes that anti-realists have nothing else convincing to say on the topics of progress and objectivity. Whether that is so may depend on the particular version of anti-realism entered into competition. But why even bother trying to develop an acceptable anti-realist alternative? Given its apparent virtues in accounting for objectivity and progress, why not sit tight with realism?

4 The realist’s problem: epistemology

Most contemporary accounts of natural kinds presume that we learn about the identity of the natural kinds, either directly or indirectly, by inspecting the categories and classifications in use in mature sciences. For this policy to be truth-conducive the scientific categories and the natural kinds must to some degree line up with one another, they must be what I will call *coordinated*. And coordination is not guaranteed; conceivably, the natural kinds might have been anything, and thus might have had no resemblance to the categories of science. If we take for granted a division of the universe into individuals, it seems that the natural kinds might have been any member of the power set of those individuals.

Of course, positively establishing coordination between our scientific categories and the natural kinds—ruling out the possibility that the universe is maniacal and our scientific categories profoundly misleading—would be a difficult, perhaps even impossible, achievement. To do so would require establishing a separate and independently certified access point to the natural kinds with which our categories could be compared. A more modest aim for a theory of natural kinds is simply to *explain* coordination, that is, to show how coordination could be, by the theory’s own lights and presuming its own truth, something other than a cosmic coincidence. To offer such an explanation is to display a theory’s internal coherence, showing in particular that there is no tension between its non-skeptical presumption that we can access the natural kinds via the scientific categories, and its claim to be the correct account of natural kinds.

In providing such an explanation, it is the realist, rather than the anti-realist, who faces the more serious difficulties. These difficulties spring from what I call the

category influence hypothesis: that the contours of scientific classifications are to some degree influenced by contingent features of scientists themselves—features over and above those constitutive of their status as scientists—in addition to being influenced by the properties of the mind-independent universe. Though there are many such routes of influence, I will focus on just one: influence stemming from the categorizer’s explanatory and predictive priorities.

The influence of priorities is most apparent when scientists working in different sub-fields differently partition the same domain of individuals into kinds. For instance, ecologists classify the domain of organisms into groups such as predator and prey, while evolutionary biologists classify the same domain into genealogically-determined species groupings. Each classification is rational given the distinct patterns each type of scientist is motivated to predict and explain (Dupré 1995; Stanford 1995). Scientific priorities also appear to influence classification in cases in which all scientists happen, as a matter of fact, to enjoy largely uniform explanatory and predictive ends. For example, Robin Hendry maintains that, though our current classification of elements based on atomic number is well motivated by an interest to account for certain material transformations, “chemical classification might have developed differently had those interests been different” (Hendry 2010, p. 138). In particular, Hendry hypothesizes that had the kinetics of chemical reactions been more important to early chemists than it actually was, then the fundamental elements might well have been based, not on atomic number alone, but on atomic weight as well. Presumably even more eccentric chemical priorities would have yielded classifications more peculiar still. For instance, had scientists cared exclusively to account for the behavior of material in centrifuges—including natural ones like our own spinning planet—a classification by density, which would crosscut our current partition, may well have emerged. Similarly, a scientific focus on patterns we see as disjunctions of two or more patterns (e.g., the-dissolving-of-substances-in-water-or-their-reactivity-with-oxygen) would have had equally curious categorical consequences; any inquirer motivated to illuminate those patterns would deploy categories difficult even to describe with languages we presently have available.

But why believe the category influence hypothesis really obtains? Besides the fact that we actually observe correlations between scientific priorities and categories, the following reasoning may be persuasive. Presuming that all individuals in our universe differ from one another in some of their features, grouping such individuals into non-singleton kinds requires the categorizer to have somehow selected between these features, foregrounding some combinations of them and back-grounding others. While such selection need not be *wholly* based on explanatory and predictive priorities, it seems that these priorities will be among the significant determinants, at least in those mature sciences that are more independent of our perceptual idiosyncrasies. More specifically, in the face of innumerable patterns displayed by a given domain of individuals, any of which might be subject to scrutiny, scientists select some patterns over others as the focus of their predictions and explanations. This selection then influences their construction of schemes of categories, as they are apt to design categories appropriate to those patterns. In the simplest case, categories are selected that group together individuals sharing whatever features are judged to underpin the behaviors the selected patterns describe.

Some readers may wonder how we should understand the metaphysical status of the features just referred to, those between which scientific interests select. Our options parallel those available in the natural kind debate. The non-skeptical realist about features will believe that a mind-independently privileged set exists to which we have access, while the anti-realist will see the set as determined in part by us, reflecting our possibly idiosyncratic tendencies to individuate features in one way rather than another (e.g., using *green* and *blue* rather than *bleen* and *grue*). Although I consider deciding between these options difficult and my position officially agnostic, given that natural kind realists will usually be feature realists as well, here I will also presume features realism. In this way my challenge to natural-kind realists does not depend on any premise they might reject.¹⁰

We are now in a position to properly state the epistemological challenge facing the natural-kind realist. It springs from the realist's difficulties responding to the following question: what is the relationship between the natural kinds and the scientific categories from which our knowledge of the natural kinds derives? Standing in the way of a satisfactory answer to this question is this. On the one hand, the realist maintains that facts about which kinds are natural and which are not are fully mind-independent, determined in no way by us. On the other hand, she suggests that we infer the identity of those kinds from the categories of science. Yet, as per the category influence hypothesis, had we possessed different explanatory and predictive priorities, the categories would likewise have differed. Therefore, the scientific categories are not mind-independent. But why have categories determined *in part by us* coordinated with natural kinds determined *in no way by us*? Without an account of this, our successful identification of the natural kinds via the scientific categories will appear to be a remarkable instance of cosmic good luck. In particular, we were lucky that we developed interests that lined up so well with the mind-independent natural kinds.¹¹

It is worth considering a few realist reactions to this coordination challenge before developing an anti-realist view that, among other virtues, avoids the difficulty. A first reply begins by observing that the category influence hypothesis eventuates in a coordination problem only on the assumption that the factors determining our scientific priorities—such as scientific curiosity and practical interests—enjoy no special truth-tracking relationship with the mind-independent metaphysical structure of our universe. Though the absence of any such relationship is reasonable as a default, the realist might undermine that assumption by describing

¹⁰ Some may wonder whether I have a right to presume realism about features on such thin dialectical ground. After all, one might worry that the kind of critique I develop in this section against natural kind realism—a form of genealogical critique—could be extended to features as well. Though I cannot explore the issue here, I believe the anti-realist is on much weaker ground when arguing against feature realism. Critiques of the kind I develop require mind-independent causal claims, causal claims that are often thought to themselves depend on facts about the features (or properties) that individuals possess. In consequence, arguably such arguments cannot be extended to the domain of features without being self-undermining.

¹¹ Though I believe it is substantial, I cannot precisely state the level of improbability ('luck') that the realist really faces. Its extent will depend on the size of the set of conceptually possible partitions of the universe into natural kinds and on how our credences are spread across them.

a mechanism by which scientists have been lead to adopt the ‘correct’ priorities, those that reflect the universe’s natural-kind structure. For instance, the realist may propose that explanatory or predictive effectiveness, the factors guiding scientists in their adoption of correct theories, will also lead them to adopt the correct priorities.

The trouble with this suggestion is that there is no reason to think that scientists animated by concerns we find unusual would do any worse by their own lights in their predictive and explanatory pursuits than we do in our own. To illustrate, consider two scientists with priorities quite different than ours. The first is fascinated by detail and aims to grasp every feature of individuals. The second, caring much less about details than do we, pays attention to only the coarsest of patterns. Both will adopt categories that seem peculiar to us. The first may appeal to singleton kinds across the high-level sciences, since these alone explain all of the behaviors he deems relevant. The second will opt for categories far less refined than our own, since our partitions respect differences that aren’t of any interest to her. Still, both of these unusual scientists will succeed as judged by their own standards. (It is not enough that these scientists will have gone wrong *by our standards*, since this would not offer a mechanism by which inquirers could detect, and eventually correct, their mistaken priorities.) The conclusion we must draw is that explanatory and predictive success cannot establish a tracking relationship between our scientific priorities and the objective natural kind structure of the universe.¹²

A second realist response is to maintain that our scientific categories coordinate with the natural kinds because the two are somehow joint products of a common cause. On this view, those features of the mind-independent universe that determined our scientific priorities (and thus the categories we adopted) also somehow set which kinds were natural.¹³ Though hardly orthodox, this proposal is properly realist in virtue of holding that the natural kinds depend exclusively on features of the mind-independent universe, and if true, it would explain how our priorities (whatever they happen to be) could reliably lead us to the mind-independent natural kinds.

This response solves the difficulty at hand, but at a price so high as to bring into relief just how deep the coordination problem is for the realist. Despite possessing no other theoretical virtues, this view depends on an extremely strong and *prima facie* implausible metaphysical assumption lacking any independent motivation: that the metaphysical structure of the universe determines our priorities in a way that will lead them to line up with the mind-independent natural kinds. To press home the peculiarity of this situation, consider how baroque the factors are that influence our scientific priorities, extending as they do from our own curiosity and practical ends to the phenomenal patterns we perceive as unitary or disjunctive. And yet this very particular combination of priorities must somehow be determined by an underlying metaphysical reality that also fixes the identity of the natural kinds *and*

¹² See Van Fraassen (1989, pp. 51–53) for a related discussion of realism about natural properties, with equally critical comments on the ability of realists to settle on the mind-independent set.

¹³ This adapts a response by Skarsaune (2011) to an epistemological objection to normative realism pressed by Street (2006), an objection that has strongly influenced the challenge to natural-kind realism presented here.

which ensures their coordination. The implication is that the deep metaphysical structure of our universe must possess a truly unfathomable complexity, going far beyond the comparatively simple realist condition that it possesses *some sort* of a natural-kind structure.

A third realist response rejects the procedure of inferring the natural kinds *directly* from the scientific categories. It was this policy that made the realist vulnerable to the coordination challenge in the first place. Instead, the realist may suggest that we access the natural kinds by consulting, not the categories of science, but the correct fundamental categorical principle. This, in concert with the non-categorical facts, will reveal which kinds are natural. In this way, perhaps we can correct for any distortions in our scientific categories caused by our particular explanatory and predictive priorities.

This reply can only evade the coordination challenge by opening it up to a closely related difficulty, the *principle-coordination challenge*. In particular, it would seem that at least some of the predictive and explanatory priorities that influence which categories scientists construct also influence which fundamental categorical principles philosophers endorse. After all, even if we do not directly infer the natural kinds from the scientific categories, naturalistically minded philosophers will agree that the latter provide a point of departure. And though broader theoretical considerations might occasionally lead us to revise our judgments about which categories form natural kinds, those judgments will remain influenced by our starting point. Moreover, those broader theoretical considerations are themselves typically influenced by our priorities, though those of a more philosophical and less scientific character. For example, philosophers both convinced that natural kinds are a mark of a science's maturity¹⁴ and committed to the maturity of particular sciences (such as biology or cognitive science) are motivated to construct accounts of natural kinds consistent with those judgments. These judgments then affect the theories of natural kinds endorsed. If, for instance, groups of heterogeneous biological individuals are to form natural kinds then natural kinds must not be defined by strict essential properties; hence, the popularity of Boyd's homeostatic property view among philosophers of biology.

If this is correct and the fundamental categorical principle a philosopher endorses is indeed influenced by her scientific and philosophical priorities, then the realist encounters the familiar difficulty of explaining the coordination between the supposedly mind-independent natural kinds and the categorical principle from which our judgments about natural kinds derive. Lacking such an explanation, it seems she must conclude their alignment is, at best, a cosmic coincidence. After all, from a conceptual point of view the correct categorical principle might have been any of an enormously large candidate pool of principles. That we settled on the correct one using the methods at our disposal thus seems to be, if not truly a miracle, at least rather unlikely.¹⁵

¹⁴ See Reydon (2009) and references therein. On the other side, Quine (1969) held that natural kinds were a mark of a science's *immaturity*.

¹⁵ The principle coordination challenge should be most worrisome for naturalistically minded realists, those who cannot reasonably deny the influence of the scientific categories on their accounts of natural

Though space has only permitted me to explore three realist responses to the coordination challenge, those responses were chosen for their promise. By their failure I fear that the realist has no very satisfying explanation of coordination to offer and thus that she must conclude we were led to the correct natural kinds only by accident. This undesirable outcome may lead the realist to the skeptical conclusion that we are much more likely wrong than right about the identity of the natural kinds. Alternatively, those wanting to avoid a skeptical conclusion might transform their metaphysical understanding of those kinds by embracing an anti-realist view.

5 An anti-realist alternative

How can the anti-realist navigate the coordination challenge? Though her explanation is simple in barest outline—that the natural kinds and the scientific categories coordinate because they are both partly determined by features of our own minds—a full explanation requires fixing on a particular anti-realist view. This, in turn, brings me to my central task: that of presenting and defending the Categorical Bottleneck Account. In constructing this view I will take as a starting point a far plainer picture: the ‘simple pragmatist’ suggestion that the natural kinds correspond to the categories that best allow us to serve our particular aims, whatever they may be.¹⁶ Although exceptionally barebones, the simple pragmatist account does pass one important test: it is able to distinguish between the actual natural and unnatural kinds in a way that at least approximates our practice of doing so. For instance, the non-naturalness of the kind “things having been transported to a distance of less than three miles from the Eiffel Tower” (Fodor 1974, p. 101) follows immediately from the fact that appealing to such a category has no use for us. Conversely, in light of the number of our projects well served by a partition of organisms into biological species—from bird-watching to predicting the spread of disease vectors to accounting for changes in the gene frequencies of populations—species taxa are deemed natural.

Still, the simple pragmatist view faces a difficulty: even if it is largely consistent with the *actual* natural kinds, it does not respect even our most confident judgments about what the natural kinds would have been in scenarios in which, while all else remained the same, our practical interests were different than they actually are. To activate your intuitions on this point, imagine that we possessed, even after concerted reflection on our total situation, very peculiar aims. With respect to animals, we might have cared only to respect the laws of kashrut, in particular the

Footnote 15 continued

kinds. What then to say to the rare metaphysician who maintains that the natural kinds have nothing to do with the categories of science? She may question my suggestion that judgments of naturalness are invariably influenced by those categories. Though I cannot engage with this opponent here, I believe a parallel genealogical critique can be constructed against the a priori principles such a non-naturalistic philosopher would need to use in constructing their purportedly category-independent theories of natural kinds.

¹⁶ Philip Kitcher articulates a view relevantly similar to this in his (2007).

division between pure animals (e.g., the cloven hoofed that chew their cud, flying insects with jointed legs above their feet) and impure animals (e.g., other flying insects that creep on all fours, water animals lacking fins and scales) described in Leviticus. Or, even more fancifully, our exclusive aim might have been to pick out members of a group constituted by three ceiling fans and a particular geranium.¹⁷ In both cases, the simple pragmatist grants just two natural kinds, in the latter case one composed of just four items and the other composed of the universe's remainder. Needless to say, this is a peculiar division, and one that even those happy to allow that *something* about us might *sometimes* affect the identity of the natural kinds will find mystifying. Whatever the character of natural kinds—whether they are sparse or promiscuous, metaphysically basic or in some way constructed—three ceiling fans and a geranium just cannot be among them.

Given that this is reason enough to abandon the simple pragmatist view, with what should it be replaced? Consider the *simple epistemic view*. According to this account, the natural kinds are groups corresponding to categories that best serve our *epistemic* aims, most centrally our project to represent aspects of the universe accurately, as well as to make correct predictions and offer up explanations about the phenomena within it that interest us. For instance, though in slightly different packaging, both Magnus (2012) and Boyd (1991) make the natural kinds relative to 'fields' or 'disciplinary matrices,' both of which can be understood as packages of epistemic aims and conceptual resources.¹⁸

Like the simple pragmatist view, the actual recommendations of the simple epistemic picture are largely in line with the natural kinds we identify. Further, it doesn't endorse outrageous conclusions about the natural kinds in cases in which our practical interests were peculiar, as they play no part in determining the natural kinds. Also speaking in its favor is the fact that this account can neatly explain coordination; it holds that our explanatory and predictive priorities are, in virtually identical ways, among the determinants of both the scientific categories and the natural kinds. This identical determination ensures that the categories actually line up with the natural kinds and also makes it the case that they would have even had our priorities been different.

A final and particularly interesting virtue of the simple epistemic view is that it yields natural kinds that are, though not fully objective, *partially objective*, meaning that its designation of some kinds as natural and others as not is invariant across assessors enjoying distinct stances or points-of-view.¹⁹ In particular, relative to a range of stances differing in explanatory or predictive aims, simple epistemic natural kinds are to some degree invariant due to a few architectural characteristics of our universe, of which I will describe two.

¹⁷ This is a peculiar aim in, among other ways, the fact that it is fully constituted by a desire to categorize things in one particular way, rather than concerning (as is far more customary) some other matter that was served by a particular categorization.

¹⁸ Note that I am overlooking many interesting properties of Magnus's and Boyd's theories, and focusing on just the feature that presents the most useful contrast with my own account.

¹⁹ Admittedly, it remains doubtful that any ordering of what is more or less partially objective could itself be objective, as that would require an objective individuation of points-of-view.

The first and most obvious characteristic underpinning the partial objectivity of the simple epistemicist's natural kinds has been emphasized by theorists going back at least to Locke: the *robust cluster* structure of our universe. This refers to the fact that the features that individuals possess are not distributed evenly across them, but rather are *clustered* together from a statistical point of view. In consequence, finding out that an individual possesses one feature often provides information about its possession of some further feature.²⁰ Though less important for this particular argument, such clusters tend to be somewhat robust in that they often would have been the same even had various aspects of the universe been different. What accounts for the existence of these clusters varies from domain to domain (e.g., high-fidelity reproduction will play a part in the explanation of robust clustering among organisms, though not among atoms). Nonetheless, these clusters always secure some stance-independence, and thus partial objectivity, for the natural kinds on the simple epistemic view. This is for the following reason. Our actual interests lead us to endorse categories that reflect certain features belonging to particular clusters. While different features may be more important for the epistemic interests of other inquirers, these features will often belong to the same clusters that we have identified. Hence, our categories will frequently coincide with theirs.

A second characteristic supporting the partial objectivity of simple epistemic natural kinds is what I call the universe's 'multi-use' architecture. This is easiest to describe with an illustration. Many biological systems, from snail ganglia to seal cerebella, possess neurons. These cells differ from one another in innumerable ways, but they share a disposition to release neurotransmitters when exposed to some neurotransmitters. This contributes to a wide variety of interesting capacities of the systems in which the neurons are embedded, and as a result, they are called upon in explanations and predictions of many different behaviors. The upshot is that the kind neuron will be relatively objective in the following sense. Our actual interests lead us to focus on systems-level behaviors that are well explained by multi-use kinds like neurons. Scientific interests different than our own may lead inquirers to different targets of explanation, but because those phenomena would also be explained by the same multi-use kinds, we will still end up with the same scientific categories, and thus the same natural kinds.

In the two ways just sketched, the simple epistemic account can avoid a highly *sensitive* dependence of the natural kinds on our own minds. This is critical because we balk at the idea that the natural kinds would have been different just because scientists focused on just slightly different phenomena. Even if the natural kinds are not totally mind independent, they ought not to be too subjective either. Yet in spite of this virtue, the simple epistemic view still faces a significant problem. The problem arises from the fact that robust clustering and multi-use architecture are not so ubiquitous in our universe as to make it the case that *any set* of properly scientific aims would lead to the same categories and kinds. In consequence, the natural kinds

²⁰ Robust clustering is most apparent when we constrain our focus to the kinds of features actually in scientific circulation and exclude highly disjunctive creations. As justifying such a constraint is notoriously difficult, I cannot pursue it here, but instead, as previously mentioned, take for granted here a broadly realist approach to 'features' in this paper.

will still depend on what our particular epistemic projects—or *disciplinary matrices* (in Boyd’s terminology)—happen to be. Consequently, the objectivity of the natural kinds on the simple epistemic view, while greater than that afforded by the pragmatist view, is still limited. For instance, though any organism will share many of its features with conspecifics, some of its still properly scientific characteristics can be shared more closely with members of other species taxa. This appears to be the case for some immunological features, which are genetically determined but highly polymorphic within recognized species, as well as for patterns in dental wear, which can be ecologically determined such that they can be shared more closely with some members of a co-located species than with some conspecifics. Had the causal consequences of these non-clustered features formed the exclusive focus of our inquiries, it seems that different categories would have been appropriate and, according to the simple epistemic view, the result would be different natural kinds.

In the light of this unfavorable implication, I recommend modifying the simple epistemic view to yield an account of natural kinds that is more robustly stance-independent and thereby more objective. According to the *Categorical Bottleneck Account* the natural kinds correspond to those categories that are metaphorical ‘bottlenecks’ in the following sense: they reflect the categories that both ourselves and a large array of scientific inquirers with epistemic aims and cognitive capacities differing from our own would sanction in common, thereby converging on a single set of categories and kinds from multiple, distinct starting positions or points-of-view. Such a view, therefore, endorses the following fundamental categorical principle: the natural kinds are groupings that match those categories that well serve actual inquirers along with (what I will call) ‘neighboring agents’—those differing somewhat from actual inquirers in their particular epistemic aims and cognitive capacities. In this way, the Categorical Bottleneck Account transforms the limited stance-independence of natural kinds from a fortuitous consequence of the architecture of our universe (as the simple epistemic account would have it) to the *defining property* of natural kinds.

As I will understand things, for a category to well serve this collection of agents does not require that it be *strictly* optimal for *all* of them, but rather that it well serve the purposes of most.²¹ The set of neighboring agents can itself be understood via an ‘epistemic agent space’ in which those located near to one another enjoy very similar aims and capacities, while those progressively farther away differ more substantially along either or both of these dimensions. The neighboring agents are those within some threshold of closeness to us. I conjecture that we can grasp the space’s structure intuitively—seeing immediately that an inquirer aiming only to elucidate patterns in the relative speed of animal locomotion is further from us than one with all of our aims, except disregarding the behavior of wombats. That we have such a space in mind and are able to reason about it should be no more mysterious

²¹ The account just described allows one to define, not just the natural kinds, but two sorts of non-natural kinds. The *pathological kinds* correspond to categories that well serve neither actual inquirers nor most neighboring agents. The *neutral kinds* correspond to categories that well serve either actual inquirers or some group of neighboring agents, but not both.

than the fact that we can grasp—though not knowing just how—that an ellipse is more like a circle than a square.²²

Why take seriously the approach just sketched? My case for it comes in three parts. First, Sect. 6 argues that the Categorical Bottleneck account can negotiate the ‘progress challenge’ from Sect. 3, and is thereby at no disadvantage to realist proposals in this regard. Second, Sect. 7 argues that it can answer the ‘coordination challenge’ pressed in Sect. 4, putting it at an epistemological advantage over realist views. Finally, Sect. 8 sketches additional respects in which the Categorical Bottleneck account is consistent with our considered judgments concerning what the natural kinds are and should be, thus speaking in its favor as against anti-realist alternatives.

6 The anti-realist response to the progress challenge

Anti-realist views across a range of subjects confront a common threat: that, by maintaining that certain facts or truths are *relative to us*, they leave no room for us to err with respect to those facts. With the possibility of error banished, progress will be equally absent. That absence speaks against anti-realist positions whenever it appears that there has in fact been progress over the history of inquiry into the subject in question. As discussed in Sect. 3, this condition is fulfilled in the case of natural kinds: it seems that we do have a better grasp on the natural kinds in (for instance) biology and chemistry now than we did prior to the Darwinian and the chemical revolutions.

Though this ‘progress argument’ may appear compelling in barest outline, I will now argue that it does not, on reflection, speak against most anti-realist accounts of natural kinds, the Categorical Bottleneck account among them. In particular, while there is one respect (examined in Sect. 7) in which the Categorical Bottleneck advocate must judge us incorrigible, there are *many other respects* along which she can make sense of error, and this is fully sufficient to make sense of the progress we detect over the history of inquiry into the natural kinds.

First, and most obviously, progress with respect to the natural kinds can be a consequence of improvements in our beliefs respecting the *non-categorical facts*, meaning all facts other than those concerning which kinds are natural and which are not (or those depending on those facts). In particular, progress in determining which features actual individuals possess is often the foundation of progress in identifying the correct natural kinds. For instance, at one time it was falsely believed that lack of reproduction in women caused a wide range of distinctive symptoms, including delusions, insomnia, and shortness of breath. Based on such beliefs, some concluded there to be a natural disease-kind of *hysteria*. We’ve since learned that living child-free has nothing like this cluster of consequences and thus discarded all scientific talk of hysteria along with the judgment that it reflects a natural disease-kind.

²² Given its role in the openly anti-realist Categorical Bottleneck account, I am able to understand both the structure of the epistemic agent space and the closeness threshold that determines the set of neighboring agents as fixed by us, rather than being properly objective.

Second, even presuming that we are right about all the non-categorical facts, we might err—perhaps through a failure of imagination or attention—in judging *natural* a kind that doesn't satisfy the fundamental categorical principle to which we are committed. Were the Categorical Bottleneck view correct this would involve judging *natural* a kind that does not correspond to a category that well serves our epistemic aims along with those of neighboring agents. For instance, we might falsely judge all morphological grouping of organisms to be natural kinds, having been struck by how many problems in the hydrodynamics of swimming and the aerodynamics of flight would be served by morphological partitions. However, even if correct in this local claim, a full survey of our interests might show that numerous of our other aims, as well as those of neighboring agents, are not advanced by such a grouping. Meanwhile, we might also find an alternative scheme of categories—an evolutionary one—that did serve those purposes well. In discovering this we might correct our error and thereby make progress with respect to the natural kinds.

Finally, we might err by misunderstanding the correct fundamental categorical principle, which for many anti-realists will mean (to a first approximation) the principle that is most consistent with our considered values and practices. This kind of error is one from which articles like this aim to free us. To illustrate, Sect. 5 suggested that it would be an error to judge the natural kinds, as the simple pragmatist does, to be groups corresponding to categories that best serve all of our actual interests, epistemic and practical. The problem was one of coherence with our practices and commitments, in particular with judgments we endorse respecting what the natural kinds would have been had our practical interests varied. Similarly, Sect. 4 suggested that it would be an error to take a strongly realist approach to natural kinds because such an approach forces us to accept an unwelcomed skepticism about the identity of those kinds. In correcting such mistakes, progress about the natural kinds—or more precisely, progress in our judgments about just *what it is to be* a natural kind—may be achieved.

Though the Categorical Bottleneck account can, in the three ways just suggested, make room for error and thus for progress, it is worth pointing out that it is not alone among anti-realist pictures in its ability to do this. Though the details will vary from case to case, the only account that straightforwardly cannot make sense of error along equivalent dimensions is the simplest style of anti-realist conventionalism, one holding that the natural kinds are those groupings that *we believe to be natural kinds*. On this view, any time our views of the natural kinds alter, the natural kinds themselves change in parallel, making progress with respect to them impossible. To avoid this threatening conventionalism, a plausible account of natural kinds must place some distance between *what we believe the natural kinds to be* and *what they really are*.

7 The anti-realist response to the coordination challenge

Even if the foregoing remarks successfully undermine the realist's progress objection, this does not mean that there is *no respect in which* the anti-realist is less

susceptible to error than is the realist. In particular, presuming (1) that we are right about the non-categorical facts, (2) that we have accurately identified the fundamental categorical principle to which our practices and norms commit us, and (3) that we have correctly applied its conditions, we cannot, according to the anti-realist, be wrong about which kinds are natural and which are not. After all, while the natural kind anti-realist may still be a realist respecting distinct subject matters (causation, laws of nature, etc.²³), she is committed to the mind-independent universe being unstructured in one specific way: in its uniting of some groups of individuals into natural kinds, and others not. In consequence, her judgments about those kinds cannot be held accountable, as the realist's kinds will be, to that structure.

How does this absence of global error affect the plausibility of the Categorical Bottleneck view? Having argued that categorical progress is possible in its absence, I will now suggest a way that the absence of global error is actually a virtue. In particular, by denying the possibility of global error some anti-realists—among them advocates of the Categorical Bottleneck account—can navigate the coordination challenge pressed in Sect. 4.

Recall, first, that the coordination challenge relied on the ‘category influence hypothesis’: the claim that the contours of the scientific categories were influenced by the projects that inquirers pursue, such that the categories (and categorical principles) would have been different had those projects been different. Given this premise, realists face a difficulty in explaining how interest-determined scientific categories (and categorical principles) came to line up with the mind-independent natural kinds. Lacking such an explanation, the realist is under pressure to judge their coordination but a cosmic coincidence. This, in turn, gives those aiming to avoid natural kind skepticism reason to abandon natural kind realism altogether. Anti-realism doesn't face a *prima facie* coordination challenge, as it doesn't claim that natural kinds have a property that the scientific categories appear to lack: mind-independence.

But, even admitting that the coordination challenge is less acute for anti-realists in general, how exactly does the Categorical Bottleneck account explain coordination? After all, the most straightforward style of anti-realist reply, the one rehearsed in Sect. 5 on behalf of the simple epistemic view, is unavailable to the Categorical Bottleneck account. That reply was *content-based* in that what allowed it to explain coordination between the scientific categories and the natural kinds was the content of the view's fundamental categorical principle. In particular, the simple epistemic account explains this coordination by citing the fact that the scientific categories and the natural kinds are both determined by the same factor—our epistemic aims. This parallel determination ensures that, even had our aims been

²³ For these subjects to be truly distinct requires that they be analyzed in a way such that they do not themselves depend on the identity of the natural kinds. I am presuming an account of causation that does this. In this respect, my view does not differ from competing views; most accounts of natural kinds in the high-level sciences presume the same with respect to causation in particular. For instance, Boyd's homeostatic property cluster theory makes facts about natural kinds depend on facts about causation—since causal claims go into defining which properties are in ‘homeostasis’—but not the reverse.

different, our judgments about the natural kinds would have nevertheless tracked the natural kinds. The Categorical Bottleneck account cannot offer a similar reply because, for the sake of achieving a greater degree of objectivity, it was designed to make the natural kinds more independent of the particular priorities of scientists.

Is epistemological obscurity the price of this more robust objectivity? Fortunately not: a *status-based* reply is still available. This reply turns on the fact that the correct fundamental categorical principle is status-dependent (as defined in Sect. 2), meaning that it is understood to be made true by us, rather than by the mind-independent universe. The coordination challenge is averted by holding that we access the natural kinds via what we take to be the correct fundamental categorical principle,²⁴ but that such principles are made true by features of us. Therefore, the Categorical Bottleneck advocate can explain coordination between the kinds we deem natural and the actual natural kinds by citing the fact that they share a common determinant: the features of us that make it the case that a particular fundamental categorical principle is correct. Those features, by hypothesis, directly determine which kinds are natural. And, presuming that we have the ability to acquire a certain kind of self-knowledge, they will also influence which fundamental categorical principle we judge to be correct, and thereby the natural kinds we infer on its basis.

But just which features of us do underpin the truth of the anti-realist's fundamental categorical principle? There are, at first glance, two broad pictures available. First, the anti-realist might maintain that the correct categorical principle *describes* the content of the rule that actually guides—with the input of non-categorical facts—the particular judgments we now make about which kinds are natural and which are not. Second, she might understand the correct fundamental categorical principle to be that which we would endorse were we to reflect on, and revise so as to achieve coherence between, all of our commitments concerning the natural kinds and related subjects.

On neither account need it be easy to uncover the correct principle, and thus success will be in no way guaranteed. Nevertheless, in contrast to the realist's obscure epistemology, the task facing us on either picture remains clear. The advocate of the first account, for instance, can maintain that we will be able to learn the correct principle by doing a kind of philosophically-informed cognitive science, reminiscent of the sort linguists undertake when they uncover the grammatical rules that guide our judgments of grammaticality, or which some moral philosophers claim to do when pursuing the principles guiding our moral judgments. The second species of anti-realist will say that we can learn the correct principle simply by uncovering our own commitments concerning the natural kinds—including those relevant to our access to them—and then, in case they are inconsistent, revising them so as to achieve consistency.

Of these two groundings for the anti-realist's fundamental categorical principle, I endorse the second 'normative' approach. Though the content of the rule that guides

²⁴ For instance, the advocate of the homeostatic property cluster theory may suggest, in light of the fact that natural kinds are groups in which clusters of properties are maintained by homeostatic means, that some particular group in which properties are so maintained is a natural kind.

our judgments about which kinds are natural and which are not is certainly *relevant* to an inquiry into the identity of the natural kinds, the springs of our present judgments cannot be the end of the story. After all, whatever the descriptive rule might be, we can always ask, having uncovered it, whether we *ought* to understand and judge natural kinds along the lines it describes; at that point it is open to us to revise our conception to suit the evaluation made. For this reason, the normative grounding, rather than the descriptive one, has ultimate authority.²⁵

8 In favor of *Categorical Bottleneck anti-realism*

The story so far: I began by distinguishing natural kind realism from anti-realism. I then considered two challenges facing accounts of natural kinds—the progress challenge (which particularly affects anti-realists) and the coordination challenge (which particularly affects realists). The Categorical Bottleneck account was then articulated and its ability to navigate these two challenges displayed.

Of course, even presuming that the progress and coordination challenges are the most critical ones facing accounts of natural kinds, further considerations not addressed here might also affect our final verdict. Or there may be other contenders that have not received due attention here. With those limitations of the present inquiry in mind, I do think the balance of reasons inclines toward an anti-realist position like the Categorical Bottleneck account. But since much of the defense of the Categorical Bottleneck account would apply to other anti-realist theories as well, I need to say something in favor of my preferred anti-realist view.

Though an account of natural kinds offering fully objective natural kinds would *prima facie* be the most attractive, I have argued that this aspiration should be abandoned in light of the epistemological difficulties that natural kind realism entails. Second-best would be natural kinds that are at least *partially objective*, where a degree of objectivity is tied to the extent to which the natural kinds would have been the same from a variety of different stances or points-of-view. In particular, it is desirable that the natural kinds be invariant across variations in the aims and the cognitive equipment of scientific inquirers.

Yet anti-realist views face difficulties realizing such partial objectivity. This is most evident for ‘success-based’ pictures, which tie the natural kinds to the categories that would allow inquirers to achieve some sort of scientific or practical success (Boyd 1990; Kitcher 2007; Magnus 2012). As we saw in Sect. 5, were hypothetical inquirers interested in investigating features not clustered with those focused on by actual scientists, working in the disciplines we have actually developed—such accounts allow that there could be different natural kinds relative to these different investigators (or to the fields in which they work).

²⁵ Of course, our normative judgments are themselves grounded in some of our actual dispositions. In this way, even on the normative approach the correctness of the correct principle depends on our actual characteristics. However, it does so in a very different way than it does for the descriptivist, and thereby recommends a different procedure for uncovering the correct principle.

Some of these accounts also run into trouble when multiple partitions of a domain would be equivalently supportive of success for some ideal inquirer, but which are not equivalently useful for us because of some feature of our innate or acquired cognitive equipment. In such cases, aim-based accounts like the simple epistemic view threaten to render *natural* merely familiar categorical schemes, those which have become easier for us to use (analogous to the way the QWERTY keyboard arrangement is superior for us based on the familiarity of contemporary typists with it).

The Categorical Bottleneck account can avoid both of these difficulties. First, the account yields natural kinds that are invariant across modifications in our particular epistemic aims. This is a consequence of the fact that a difference in our aims will not usually affect which groupings well serve *a range of aims or capacities*, which is the standard that determines whether a kind is natural on the Categorical Bottleneck view. This enables the Categorical Bottleneck account to downplay the relevance of our particular interests and to privilege the importance of aspects of the universe's causal structure like robust clustering and multi-use architectures.

Second, the account will not elevate a grouping to natural-kind status simply in virtue of our familiarity with it. Take a case in which the initial partition really could have been made differently but as effectively—analogous to the way that we might have differently divided the night sky into constellations that were just as well suited to navigational tasks.²⁶ There are bound to exist inquirers nearby us in epistemic agent space who developed initial divisions different from our own, and who are best served by the partitions that *they* are familiar with. In consequence, the divisions we actually constructed will not be strongly favored by both ourselves and our neighbors, and thus will not reflect natural kinds on the Categorical Bottleneck account.

Finally, the Categorical Bottleneck account has a distinct advantage over anti-realist competitors when we consider the role of natural kinds in the scientific enterprise over time. Throughout my discussion I have been emphasizing that natural kinds ought to be relatively stable across inquirers inhabiting epistemic agent space who have somewhat different scientific priorities than we do. However, many of these neighboring inquirers exist, not in adjacent possible worlds, but in our own future. I submit that it is particularly desirable that the categories we identify as natural kinds have some prospect of longevity in the face of interests and priorities that will almost certainly evolve time. There are, of course, some possible changes in our interests that would leave Categorical Bottleneck categories without a scientific calling. Still, a focus on categories satisfying the fundamental categorical principle that the Categorical Bottleneck account recommends will give us a good chance of developing categories that will indeed prove to be a valuable cultural inheritance. Categorical schemes, after all, are labor-intensive achievements, and it is best if our labor goes into schemes that will suit future inquirers as well as ourselves.

²⁶ See Magnus (2012) for a helpful discussion of this sort of situation, though approached in a different way than I do here.

But what if you disagree with this final claim, that the natural kinds ought to provide a ‘categorical legacy’ to future generations? What if, instead, you believe we ought to prioritize those categories that—to the degree possible—bias future generations towards investigating some kinds of questions over others? If so, you might endorse a different account of natural kinds—as in your case a different fundamental categorical principle may well best cohere with the rest of your commitments. In this way, even if the Categorical Bottleneck natural kinds do not depend on our particular epistemic interests and priorities, they still ultimately depend on our values, particularly on what purposes we think the entire theoretical apparatus of science should serve.

Acknowledgments For helpful discussions of this paper, thanks to Richard Boyd, Marc Ereshefsky, Andrew Franklin-Hall, Sharon Street, Michael Strevens, participants in the Corridor Reading Group (Errol Lord, Barry Maguire, John Morrison and Kristin Primus), and an audience at Cornell University. Many thanks are also due to Elizabeth Radcliffe for organizing this issue and for her patience with my contribution to it.

References

- Boyd, R. (1990). Realism, conventionality, and ‘realism about’. *Festschrift for Hilary Putnam*, pp. 171–195.
- Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies*, 61(1), 127–148.
- Boyd, R. (1999a). Kinds, complexity and multiple realization. *Philosophical Studies*, 95(1), 67–98.
- Boyd, R. (1999b). Homeostasis, species, and higher taxa. In R. Wilson (Ed.), *Species: New interdisciplinary essays*. Cambridge, MA: The MIT Press.
- Boyd, R. (2000). Kinds as the “workmanship of men”: Realism, constructivism, and natural kinds. *Rationalität, Realismus, Revision: Vorträge des*, 3, 52–89.
- Boyd, R. (2010). Realism, natural kinds, and philosophical methods.
- Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.
- Darwin, C. (1859). *On the origin of species, or the preservation of favoured races in the struggle for life*. London: J. Murray.
- Dupré, J. (1995). *The disorder of things: Metaphysical foundations of the disunity of science*. Cambridge, MA: Harvard Univ Pr.
- Elder, C. (2004). *Real natures and familiar objects*. Cambridge, MA: MIT Press.
- Ellis, B. D. (2001). *Scientific essentialism*. Cambridge, MA: Cambridge Univ Pr.
- Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28(2), 97–115.
- Ghiselin, M. T. (1997). *Metaphysics and the origin of species*. Albany, NY: Suny Press.
- Goodman, N. (1978). *Ways of worldmaking*. Indianapolis: Hackett Publishing Company.
- Hacking, I. (2007). Natural kinds: Rosy dawn, scholastic twilight. *Royal Institute of Philosophy Supplement*, 61(1), 203–239.
- Hendry, R. F. (2010). The elements and conceptual change. In *The Semantics and Metaphysics of Natural Kinds*, pp. 137–158.
- Hull, D. L. (1978). A matter of individuality. *Philosophy of Science*, 45, 335–360.
- Hunter, E. J. (2009). *Classification made simple: An introduction to knowledge organisation and information retrieval*. Farnham: Ashgate Publishing.
- Khalidi, M. A. (1998). Natural kinds and crosscutting categories. *Journal of Philosophy*, 95(1), 33–50.
- Kitcher, P. (1984). Species. *Philosophy of Science*, 51, 308–333.
- Kitcher, P. (2001). *Science, truth, and democracy*. New York: Oxford University Press on Demand.
- Kitcher, P. (2007). Does ‘Race’ have a future? *Philosophy & Public Affairs*, 35(4), 293–317.
- Kitcher, P. (2008). Carnap and the caterpillar. *Philosophical Topics*, 36(1), 111.

- Kornblith, H. (1995). *Inductive inference and its natural ground: An essay in naturalistic philosophy*. Cambridge, MA: The MIT Press.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103(412), 473–490.
- MacLeod, M., & Reydon, T. A. C. (2012). Natural kinds in philosophy and in the life sciences: Scholastic twilight or new dawn? *Biological Theory*, 1–11.
- Magnus, P. D. (2012). *Scientific inquiry and natural kinds*. New York: Palgrave Macmillan.
- Mill, J. S. (1882). *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*. Green: Longmans.
- Psillos, S. (1999). *Scientific realism: How science tracks truth*. Psychology Press.
- Quine, W. V. (1969). Natural kinds (pp. 5–23). Netherlands: Springer.
- Reydon, T. A. (2009). Do the life sciences need natural kinds? *Croatian Journal of Philosophy*, 26, 167–190.
- Scerri, E. R. (2011). *The periodic table: A very short introduction*. Oxford: OUP.
- Sider, T. (2012). *Writing the book of the world*. Oxford: Oxford Univ Pr.
- Skarsaune, K. O. (2011). Darwin and moral realism: survival of the fittest. *Philosophical Studies*, 152(2), 229–243.
- Stanford, P. K. (1995). For pluralism and against realism about species. *Philosophy of Science*, 65, 70–91.
- Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109–166.
- Street, S. (2008). Reply to Copp: Naturalism, normativity, and the varieties of realism worth worrying about. *Philosophical Issues*, 18(1), 207–228.
- Street, S. (ms.). Normativity and water: The analogy and its limits.
- Van Fraassen, B. C. (1989). Laws and symmetry.