Brit. J. Phil. Sci. **0** (2019), 1–19

Can the Epistemic Value of Natural Kinds Be Explained Independently of Their Metaphysics?

Catherine Kendig and John Grey

ABSTRACT

The account of natural kinds as stable property clusters is premised on the possibility of separating the epistemic value of natural kinds from their underlying metaphysics. On that account, (i) the co-instantiation of any sub-cluster of the properties associated with a given natural kind raises the probability of the co-instantiation of the rest, and (ii) this clustering of property instantiation is invariant under all relevant counterfactual perturbations. We argue that it is not possible to evaluate the stability of a cluster of properties without taking stock of the metaphysical picture used to account for that stability. Thus, even on the stable property cluster account, the epistemic value of natural kinds remains partly grounded in their metaphysical status.

1 *Introduction*
2 *Cliquish Stability and Natural Kindness*
3 *Errors about Cliquish Stability: Three Cases*
  3.1 *Race and IQ*
  3.2 *Lichen symbionts*
  3.3 *Man o' war as jellyfish*
4 *Epistemic Value and Metaphysical Presuppositions*
5 *Conclusion*

1 Introduction

In a recent but already influential article, Slater ([2015]) suggests that we shift the investigation of natural kinds away from inquiry into the metaphysics that underwrite such kinds. Instead, he proposes to examine those features of natural kinds that make legitimate their use in scientific inquiry. To this end, he develops an account of natural kindness, which is 'a kind of status things can have that partially underpins their role in our inferential practices'

doi:10.1093/bjps/axz004

(Slater [2015], p. 378). Slater's plan is thus to set aside questions about whether we can give a uniform account of the processes or mechanisms (or what have you) that account for the clustering of properties that characterize natural kinds. On his more modest approach, all that we require is an account of what it is about such property clusters that allows us to make cogent inferences from the instantiation of some properties in the cluster to the probable instantiation of the rest.

What is the relevant feature of these clusters of properties that allow us to draw such inferences? In a word: stability. Slater ([2015], p. 396) 'requires only that these properties be sufficiently stably co-instantiated to accommodate the inferential and explanatory uses to which particular sciences put [natural kinds]'. When we encounter an organism with a large trunk, roots, and broad leaves, for instance, it is the stability of the co-instantiation of properties like 'having broad leaves and being capable of photosynthesis' that underwrites our inference that this organism probably photosynthesizes.

Much of Slater's account is thus dedicated to spelling out the meaning of 'sufficiently stable' co-instantiation of properties. He presupposes that it is possible to work out these epistemological details without saying anything about what grounds such stability—that is, without saying anything about the homeostatic mechanisms or processes that explain the fact that the presence of one property is likely to be accompanied by the presence of certain others. It is not that the inquiry into the grounds of stable clustering is useless; it is simply an entirely distinct inquiry than Slater thinks is required for explaining the epistemological significance of natural kinds.[1] As he puts it, 'We can grant that the project of uncovering certain homeostatic mechanisms underlying the stability of some properties can contribute to the construction of epistemically fruitful classification schemes without supposing that such identification is necessary or that the mechanisms should be part of the philosophical analysis of natural kinds' (Slater [2015], pp. 402–3).

Slater thus holds what we will call the 'independence of epistemic value of natural kinds': the epistemic value of natural kinds can be understood independently of their metaphysical grounds. Surely there is something right in this claim; Slater's analysis largely succeeds in providing an analysis of natural kinds that is metaphysically neutral. Yet there are reasons to think that the epistemic independence of natural kinds has its limits. Although Slater's general account presents natural kinds with careful neutrality as to their metaphysics, we will argue that it is not possible to accept the epistemic fruit of any particular natural kind without some examination of the metaphysical roots that lie at its base. As such, we offer our analysis as a friendly amendment to

[1] As such, Slater ([2015]) conceives of his project as orthogonal to those that are focused on inquiries into the nature of natural kinds, often cashed out in terms of their underlying homeostatic properties or mechanisms (Boyd [1991], [1999], [2000], [2010]).

Slater's approach and a possible addendum to others who seek a metaphysically neutral account of natural kinds. Notably, while some recent critics of approaches such as Slater's—including Lemeire ([2018]) and Martinez ([2017])—have argued that epistemology-only accounts of natural kinds are doomed to failure, we are more optimistic. Our aim is only to highlight a price that an account like Slater's must pay: in order to commit to the independence of epistemic value of natural kinds, one must leave all metaphysical theorizing entirely up to the scientists who make use of the kind in question.

We suggest that purportedly epistemology-only accounts like Slater's provide a helpful account of the epistemic value of natural kindhood when the attribution of kindhood to cliquish property clusters gets it right, but successful attributions of kindhood depend on accurate metaphysical presuppositions that a metaphysically neutral account cannot address. Scientists might not call what they are doing 'metaphysical theorizing' but that doesn't mean they aren't doing it (even if they only do so for pragmatic reasons, such as to enable counting or measuring).

Section 2 sets out Slater's account of the stable property clustering involved in natural kinds. On that account, (i) the co-instantiation of any sub-cluster of the properties associated with a given natural kind raises the probability of the co-instantiation of the rest, and (ii) this clustering of property instantiation is invariant under all relevant counterfactual perturbations. Sections 3 and 4 then present our argument against the full-fledged independence of epistemic value of natural kinds. The core of our argument is that it is not possible to make good judgements about the invariance or stability of a cluster of properties (condition (ii) in Slater's account) without taking stock of the metaphysical picture used to account for that stability. Finally, Section 5 concludes.

2 Cliquish Stability and Natural Kindness

Slater calls the requisite stability of co-instantiation 'cliquish stability'. A property cluster $[\Phi]$ is cliquishly stable under two conditions:

> Cliquishness: For any individual x, the fact that x instantiates any sub-cluster of properties in $[\Phi]$ makes it probable that x instantiates all of them.
> Stability: The cliquishness of $[\Phi]$ is invariant under all relevant counterfactual perturbations.[2]

[2] Slater explicitly develops his account of cliquish stability by analogy with Lange's ([2009]) account of the stability of natural laws in terms of their invariance under certain counterfactual perturbations. We are paraphrasing Slater's account, which is quite brief. Slater ([2015], p. 400) lays out the view as follows:

> [. . .] a property cluster $[\Phi]$ is cliquishly stable when for all x and for many sub-clusters $[\Phi_1]$, $[\Phi_2]$, $[\Phi_3]$, . . .:

For example, suppose that the cluster of properties associated with a putative kind tree is cliquishly stable; what does this amount to? Well, suppose you encounter a new organism that looks like a tree—it has a trunk, branches, leaves, and flowers. You recognize, in other words, that it has a sub-cluster of the properties that characterize the kind tree. However, the cluster of properties characterizing the kind 'tree' includes many features that you are not presently able to investigate directly: [being deciduous, possessing a xylem, phloem, and cambium, capable of bearing fruit (if female or hermaphrodite)]. Nevertheless, given that this cluster of properties is cliquish, you judge reasonably when you judge (say) that the leafy thing with roots is probably an angiosperm (a flowering plant). The fact that it has a sub-cluster of properties associated with the kind 'tree' raises the probability that it has all of them.

But suppose you start to wonder: how secure is this sort of probabilistic inference? This is where stability comes into the picture. The fact that the cluster of properties characterizing the putative kind 'tree' is stable implies that there are no relevant ways the world could be (or could have been) such that this organism would possess the properties you have observed, but fail to (probably) be a tree. The changes that would be required are either too radical or too irrelevant to undermine the probabilistic entailments at issue. There are ways that the world could be (or could have been) such that the probabilistic entailments involved in cliquishness would fail. But we can ignore them:

(i) When evaluating stability, we ignore counterfactuals that are inconsistent with the cliquishness of the given cluster: Suppose fruit trees had evolved to have different properties than they actually have. Then, of course, the probabilistic entailment from a sub-cluster of fruit tree properties to all the rest would fail. But this is not relevant to the question of the inferential role of natural kinds.

(ii) When evaluating stability, we ignore counterfactuals that are inconsistent with natural law: Suppose the natural laws had been different, such that (say) flowers weren't required for fruit trees to bear fruit. Then it might not be reasonable to infer that the fruit-bearing leafy organism

$$p \mathrel{\Box\!\!\rightarrow} (([\Phi_1]x \Rightarrow [\Phi]x) \wedge ([\Phi_2]x \Rightarrow [\Phi]x) \wedge ([\Phi_3]x \Rightarrow [\Phi]x) \wedge \ldots,$$

$$q \mathrel{\Box\!\!\rightarrow} (p \mathrel{\Box\!\!\rightarrow} (([\Phi_1]x \Rightarrow [\Phi]x) \wedge ([\Phi_2]x \Rightarrow [\Phi]x) \wedge ([\Phi_3]x \Rightarrow [\Phi]x) \wedge \ldots,$$

$$r \mathrel{\Box\!\!\rightarrow} (q \mathrel{\Box\!\!\rightarrow} (p \mathrel{\Box\!\!\rightarrow} (([\Phi_1]x \Rightarrow [\Phi]x) \wedge ([\Phi_2]x \Rightarrow [\Phi]x) \wedge ([\Phi_3]x \Rightarrow [\Phi]x) \wedge \ldots,$$

where p, q, r, ... meet the following conditions:

(a) they are consistent with the probabilistic entailment relationships from sub-clusters to clusters;

(b) they are consistent with the natural laws...;

(c) they meet the relevant applicability standards.

with a woody trunk before you is an angiosperm. But such counterlegal considerations do not seem relevant to the inferential role of natural kinds in our actual world.

(iii) When evaluating stability, we ignore counterfactuals that are irrelevant: Suppose you lived near someone who very much liked to make robots that were lifelike replicas of trees. Then the fact that an individual appeared treelike would not make it probable that it would also have the other properties in the cluster: it's as likely to be a robot as a tree. But, again, such arcane worries should not in general undermine our confidence in the cliquish stability of a cluster of properties, and so we ignore them.

In his own discussion of these matters, Slater recognizes that the third category of counterfactuals to be ignored when evaluating stability is the most problematic. We agree. Slater's concern is over the troubles that might arise if our judgements about which counterfactuals are relevant—the third category above—relies on scientists' interests, so that they (rather than the world) are the real adjudicators of when property clustering is stable (see also Craver [2009]). To preserve 'naturalness' if not realism about natural kinds, Slater suggests that we swap disciplinary 'interests' with disciplinary 'relevance'. If we do this, we can then tie cliquish stability to a particular scientific context, theory, or project (Slater [2015], p. 402). This context-driven approach to defining cliquish stability has the result that some property clusters may only be considered natural kinds within particular disciplines or scientific research projects.

This seems to lead to the conclusion that the status of natural kindness may be determined only for a pre-selected domain of inquiry. It seems that rather than seeking to ground mechanisms of homeostatic properties, he prefers instead to put 'stability at the ground level of an account of natural kinds' (Slater [2015], p. 403) in hopes of a 'metaphysical[ly] neutral' position. By sidestepping the metaphysical questions, Slater aims to focus on a discipline-specific view of natural kinds. The relevance of various counterfactual perturbations is determined within the different disciplines. But this is surely also where (if there are any) the underlying properties on which homeostatic mechanisms depend may also be revealed. Slater does not deny the importance of attempts to reveal these properties, but sees them as routes scientific research may take rather than a fruitful subject for philosophical inquiry. The role of relevance seems to be to circumscribe kinds in relation to either potential or actual domains of inquiry (Slater [2015], p. 404).

To summarize, Slater posits a meta-level analysis of natural kinds in terms of the status of natural kindness. On his view, providing a metaphysical explanation of why certain property clusters are afforded this status makes no

further contribution to our epistemology of natural kinds. One difficulty this view raises is that, if natural kindness is construed as stable, cliquish property clustering, but we have no account of the mechanisms underlying this clustering, it is unclear how to decide which property clusters we should be tracking as potentially stable. To see how this difficulty manifests, we next consider some cases in which natural kind attributions are mistaken. Such errors serve to highlight the limitations of Slater's metaphysically neutral approach.

3 Errors about Cliquish Stability: Three Cases

Given Slater's analysis, there are two main ways attributions of stability—and thus judgements about natural kindness—might go wrong. The first type of error is to attribute stability to a cliquish property cluster [Φ] that in fact lacks it; call these 'false positives'. The second type of error is to deny that [Φ] is stable when in fact it is; call these 'false negatives'.

False positives occur when, upon observing a cluster of properties that appears stably cliquish, scientists neglect to consider some relevant counterfactual perturbations that would reveal the cluster to be unstable. That is, such errors arise from ignoring relevant counterfactuals. In cases of this sort, there is a property cluster that is actually cliquish, such that (say) repeated experiment indicates that the co-instantiation of some properties in the cluster makes it probable the others are co-instantiated as well. Yet there is a relevant possible world where this cliquishness fails, so that the cluster lacks the appropriate coherence to qualify as a natural kind after all.

False negatives, by contrast, arise from giving weight to irrelevant counterfactuals. In such cases, there is a property cluster that is actually cliquish, but which is mistakenly thought to be unstable. That is, it is thought that there are relevant possible worlds at which the cliquish co-instantiation of the cluster fails. However, the worlds taken to falsify the cluster's stability turn out to be irrelevant. If all of the irrelevant counterfactual perturbations were set aside, then the property cluster in question would be recognized as stable. But with the irrelevant counterfactuals in the mix, the cliquishness of the cluster is thought—wrongly—to be unstable in ways that undermine its claim to natural kindness.

The counterfactuals and possible worlds at issue should not be seen as mysterious things. It is helpful to think loosely of these possible worlds as the ways things would turn out if, holding fixed nature's laws and history, we undertook an experiment that has not been (and might not ever be) actually performed. In some of these scenarios, the experiment might show us that a property cluster we took to be a natural kind is not one. In other scenarios, the experiment might show the reverse. A cluster that we thought to be only accidentally cliquish might surprise us, retaining its cliquishness in spite of

our experimental intervention. What is important about these scenarios is that, on Slater's account, the epistemic value of natural kinds depends on our judgements about the relevance or irrelevance of such possibilities.

Notably, both of these types of error are likely to arise if one has a mistaken explanation of the stability of the given cluster of properties. For both types of error can easily result from misidentifying the mechanisms responsible for the stability that explains why cliquish clusters are cliquish. Good explanations of natural kindness rely not just on cliquishness as an observable and statistically measurable property of clusters, but on the hypothetical claim that these clusters' cliquishness is (and will be) stably maintained under various counterfactual changes. Scientists relying on a false account of the mechanism or ground of the property cluster are vulnerable to considering the wrong counterfactual changes. Remedying such errors in such cases requires correcting the explanation in question. Some examples of each type of error should illustrate the general point.

3.1 Race and IQ

One example of an instance where false positives have occurred is in the attribution of natural kindness to the relationship between various morphological characteristics that have been used to determine a person's race and their IQ. Taken as a cluster of properties that appears stably cliquish, Herrnstein and Murray ([1994]) linked racialized morphological characteristics and social behaviours with an individual's IQ test results.[3] Doing so is possible only by neglecting relevant counterfactual perturbations that would reveal the cluster of racialized characteristics and the results of IQ tests to be an unstable grouping.

Although there are a host of epistemic and value errors that lead to the erroneous assertion that racialized morphological characteristics and IQ constitute a group of properties that is cliquishly stable and deserving of the status of natural kindness, we focus on the metaphysical errors this mistake involves. The erroneous identification of this cluster of properties as cliquishly stable arises because certain counterfactual scenarios are ignored. To see this, we need to consider cliquishness and stability as two different conditions. For instance, racialized morphological characteristics and IQ are cliquish in studies that already assume a racialized connection between IQ and race. In these research contexts, certain racialized traits and the results of IQ tests would indeed turn out to be cliquish. But what leads some researchers to say these cliquish properties are stable is that they are presupposing a false (racist)

[3] Despite widespread and sustained criticism since its publication, *The Bell Curve* has, regrettably, had a recent surge of sales (Siegel [2017]).

picture of what grounds their cliquishness. Specifically, this conclusion depends on the erroneous presupposition that race and intelligence are both based on a set of shared genes that are reliably inherited by all (or most) people who can be morphologically identified (either by a clinician determining the race of a research subject on the basis of a set of racialized traits, or by the self-identification of a person's race by the participants themselves).

The move from identifying this cluster of properties as cliquish to the suggestion that it is stable thus relies on an underlying metaphysical grounding of race and intelligence. That this mistake can occur, means that scientists may wrongly attribute the status of natural kindness to certain putatively cliquish but unstable property clusters because they are using an erroneous metaphysics. What happens in the case of race and IQ (most egregiously in the published research of Herrnstein and Murray in *The Bell Curve*) is that natural kindness is attributed in a place that it shouldn't be.

3.2 Lichen symbionts

Another instance of false positives might occur if we fail to attribute natural kindness correctly because we have failed to accurately recognize the mechanisms of stability. The second example we discuss focuses on the metaphysics of the lichen symbiont. It illustrates what happens if we fail to recognize the appropriate mechanism of stability necessary for the lichen to grow.

Lichens have long been studied and defined as a two-part mutualistic symbiont composed of a fungus and a photosynthetic partner like an algae or cyanobacteria (Nash [2010]). The lichen was understood to be composed of a fungal partner (the mycobiont) that was thought to be able to produce a lichen in collaboration with a photosynthesizing algal partner (the photobiont) (Nash [2010]). The standard view has been that lichens are systems that have one fungus—typically an Ascomycete or Basidiomycete. Although other fungi are known to be parts of the lichen system in a less functional or evolutionarily impactful role as parasites, the classical binary view of lichen composition of mycobiont-photobiont has been widely accepted (Nash [2010]). The cluster of properties characterizing the kind 'lichen' includes: the capacity to form microfilamentous, microglobose, or crustose thalli where photobiont cells are kept, sustained, and controlled (Honeggar [2010], p. 36); being poikilohydric[4] rather than homiohydric (like most flowering plants are) (Green *et al.* [2010], p. 165); and exhibiting perennial growth.[5]

[4] *Poikilohydric* organisms are organisms whose water status passively changes depending on the availability of water in the local environment, (for example, in the form of rain, dew, fog, or humidity) (Green *et al.* [2010], p. 165). Lichens in harsh Antarctic conditions can rehydrate after eight months of being dry and are then capable of photosynthesis shortly thereafter (Kappen *et al.* [1998]).

[5] In some cases, this perennial growth can extend past 1000 years (Beschel [1961]; Pringle [2017]).

According to the standard view, the explanation of stability of this cluster of properties was supposed to be the presence of the same mycobiont over time—although a photobiont must be present in order to form the lichen, the species of photobiont may not even be identified. This metaphysical view so dominates lichenology that it underpins classificatory practices that rely on a single fungus (not the algae or cyanobacteria or any secondary fungus that might be part of the symbiotic system) to name and track lichens. It is the single lichenizing fungus together with the photobiont that has been thought to ground the stability of the properties characteristic of the lichen symbiont.

However, this bipartite metaphysical picture of the lichen has been challenged. Recent research suggests that this one-lichen, one-fungus metaphysics ignores relevant counterfactual alternatives: the existence of a lichen may require three or more (rather than two) organisms (Rambold and Triebel [1992]; Lawrey and Diederich [2003]; Henskens *et al.* [2012]; Chagnon *et al.* [2016]; Spribille *et al.* [2016]). The recent discovery that some lichens are made up of three or more rather than two symbiotic parts (as was widely thought) employs a revised metaphysical notion of the lichen correcting (at least in certain species) the use of an erroneous metaphysical assumption. Lichens are not always bipartite. They may be composed of one mycobiont and (not one but) two photobionts. For instance, cyanolichens are not bipartite but tripartite, possessing both an algal as well as a cyanobacterial partner (Henskens *et al.* [2012]). The symbiotic lichen system of mycobiont and photobiont may possess other microbial species (for example, species in the genera *Hypoxylon* and *Daldinia*) or it may have a third fungal partner that helps maintain the shape and structure of the lichen thallus (U'Ren *et al.* [2016]). For example, Chagnon *et al.* ([2016]) found that endolichenic microbial species present within the lichen thallus play a significant role in the way lichens assemble, in the causes of their formation, and in making the instantiation of the associated properties stable over time. Previous studies ignored some counterfactual perturbations relevant for lichenization. They ignored many constituents of the cluster of properties of lichen that are responsible for its being stably instantiated. Stability was mistakenly thought to be due to one fungus and one photobiont. The case is a false positive because the cliquishly stable properties are not explained by its one-lichen, one-fungus, bipartite composition. In many lichens, stability relies on the presence of not two but three symbiotic microbial species.

The two examples of false positives highlight problems with relying on characterizations of cliquishness alone when attempting to make kind attributions. Accurate kind attributions depend on accurate metaphysical presuppositions, which are themselves tacit in the putative explanations of cliquish stability we adopt. While Slater's account of natural kindness is offered as a metaphysically neutral alternative to that of Boyd's homeostatic property cluster theory (Boyd [1991], [1999], [2000], [2010]), adopting it means we

still need to consider cases where views about the metaphysical grounds of property clustering have resulted in errors of natural kind attribution. The account of natural kinds in terms of cliquish stability has at least two undesirable limitations. Following the first case, stability is liable to be misattributed if the mechanism of stability is erroneously determined. Secondly, cliquish stability cannot be used to determine whether or not the clusters themselves are natural. Any kind attribution that relies on these clusters is anchored to a particular metaphysical picture of the world that it takes to be natural without being able to arbitrate between different candidate pictures of the world.

We now describe an instance of a false negative—that is, a case where attention is given to irrelevant counterfactuals as evidence that the subject of investigation does not constitute a stable set of properties. In doing so, the cluster in question is determined to not be stable and is thereby not attributed the property of natural kindness.

3.3 Man o' war as jellyfish

Portuguese man o' war (*Physalia physalis*) have long been thought of as a kind of jellyfish as their outward morphological structures of a float and tentacles appear similar to those possessed by jellyfish. Although man o' war belong to the same phylum, Cnideria, they are a species of siphonophore that develop very differently from jellyfish. Man o' war are floating hydrozoans that are composed of four types of animals. The pneumatophore, a gas-filled polyp, makes up the distinctive clear blue float. The dactylozooids are a second type of organism that make up the stinging cnidocyte-covered tentacles that are used for defence and to paralyse and kill prey (Kurlansky [2002]). Those organisms that ingest food are the gastrozoids, and the organisms that the man o' war relies on for reproduction are the gonozoids (Kirkpatrick and Pugh [1984]). Each of these zooids has their own nervous system. The man o' war does not have a central nervous system and its nervous system is unlike the neural net of the jellyfish. Because of this, the tentacles of a man o' war are free swimming and so it is unable to coordinate behaviour. Man o' war can float but cannot propel themselves. By contrast, true jellyfish, like Scyphozoa, Cuboza, and Stauroza, are single organisms that are capable of propulsion by expanding and contracting their bodies to push water through themselves in order to move.

Jellyfish have coordinated nervous systems structured as a nerve net. The nerve net is not a central nervous system but has the ability to transmit impulses from one part of the organism (the tentacles) to another part of the jellyfish's body. The nerve net, the circular nerve ring, and the rhopalial lappet allow the communication of stimuli from one part of the body to another via its nerve cells (Bardi and Marques [2007]). If the features of a jellyfish—such as the capacity for propulsion, possession of a coordinated nervous system

structured as a nerve net, ability to sting prey for consumption, and possession of tentacles—were taken as the model against which to decide whether or not a man o' war's properties were cliquishly stable, then the man o' war's properties—possessing tentacles to sting prey for consumption, possession of a float, but unable to propel itself—would fail to count as being cliquishly stable. This is because they were not the result of a coordinated nerve net. Yet that judgement would be a mistake. The determination that the man o' war's properties are not cliquishly stable would be the result of attending to irrelevant counterfactual perturbations—if the man o' war had a nerve net then its properties would be cliquishly stable. Even though it might be the case that all parties would recognize that there is a certain biological individual at issue—the man o' war—someone using the model of a jellyfish and looking at the list of properties a man o' war actually has would deny that those properties constitute a natural kind. The properties of the man o' war are cliquishly stable, but their stability is not grounded in the possession of a unified nerve net possessed by a single organism (like jellyfish). Instead, the man o' war's properties are the result of a colony of functionally specialized organisms working together. Recognizing this alternative metaphysical model of the unity of the man o' war allows us to see that its properties constitute a natural kind only by understanding them in light of a different metaphysical model.

It may be illustrative to compare this error to a more familiar sort of mistake. For example, we might mistakenly believe the stability of cliquish traits in whales is caused by the same mechanism of stability that is present in fish. A consequence of this false analogy may be that we expect whales to have gills instead of lungs.[6] The false belief that whales have gills would be due to our having failed to make an observation, perform a dissection, or conduct an experiment that is possible (and perhaps even easy) for us to do or has actually been done by someone else. The whale case may initially appear to be exactly similar to our man o' war case. Shedding light on how the man o' war case and the whale case are different may help to highlight the motivation for our specific concerns in this section. Both the whale and the man o' war case are instances of a false negative. But, though both are false negatives, the whale case is not a particularly worrying instance of a false negative. In Slater's terms, the relevant property cluster would be the cluster consisting of the properties characteristic of the kind 'whale' but including the property of having gills in place of having lungs. However, there are counterfactual perturbations that would show this cluster to be unstable. What makes the whale case seem innocuous is that those possible perturbations are relatively easy to access: we can just do the experiment.

[6] Thanks to an anonymous referee for suggesting this as a potential parallel case.

What the case of the man o' war is intended to show, by contrast, is that there are other instances where the relevant mechanisms of stability are not so easy to recognize or test. These are the instances that we are especially interested in—the cases where scientists are making metaphysical assumptions that not only outstrip empirical evidence, but where the means by which these assumptions can be tested may not even lie in nearby possible worlds. That is, false positives and false negatives can still exist in cases where scientists have ready access to data that would suggest a modification of their metaphysical view. But they can also exist in cases where scientists do not, or do not yet, have access to data that would suggest revision to their metaphysical picture.

4 Epistemic Value and Metaphysical Presuppositions

Where does all of this leave us? Recall that, on Slater's account, to ask whether a cluster of properties constitutes a natural kind is to ask whether those properties are cliquish (that is, the instantiation of any sub-cluster raises the probability of the instantiation of the rest), and whether their cliquishness is stable (that is, it survives all relevant counterfactual perturbations). These features are what make natural kinds epistemically valuable. Slater contends that neither of these features requires us to make any claim about the underlying metaphysics of natural kinds. Thus, given that cliquish stability is an adequate characterization of the epistemic value of natural kinds—a point we do not contest—it follows that their epistemic value can be understood independently of their metaphysical ground.

With that in mind, the upshot of the cases we have discussed is this. Attributions of stability (or instability) are always made on the basis of assumptions about which counterfactual perturbations are relevant. Those assumptions include certain underlying metaphysical commitments. The epistemic value of a natural kind is thus contingent upon those metaphysical commitments. This severely limits the extent to which we can understand the epistemic value of natural kinds independently of considerations pertaining to the metaphysics of natural kinds.

More carefully, Slater's claim that the epistemic value of natural kinds can be understood independently of any metaphysical account of natural kinds could be understood in several different ways. On some (very broad) readings, we claim it will be false. Only on other, more carefully restricted, readings is it supported by Slater's analysis. Consider first:

> Strong Independence: For any purported natural kind K, the epistemic value of K can be fully captured by giving a metaphysically neutral account of its inferential role.

This strong version of the independence thesis is false. The problem is that if we are considering the domain of all purported natural kinds, the epistemic value of any given kind will be contingent upon the success or failure of our attributions of cliquish stability. And those attributions, we have argued, frequently hinge on the success or failure of the metaphysical model that is used to make decisions about the relevance or irrelevance of certain counterfactuals.

However, there is a weaker version of the claim that survives:

> Weak Independence: For any natural kind K, the epistemic value of K can be fully captured by giving a metaphysically neutral account of its inferential role.

If we know that K is a natural kind—that is, we know that the cluster of properties associated with K is cliquishly stable—our situation is much improved. At that point, if someone wishes to challenge the utility or epistemic value of our appeals to K, we can simply appeal to the role that K (successfully!) plays in our inferences. The metaphysical background against which we determined that K is a natural kind now drops out of the picture. Yet this is not to say that the epistemic value of K is independent of that metaphysical background. It is only that, once the relevant metaphysical assumptions are taken for granted (and reflected in the fact that K is stable), there is no need to mention them explicitly in our account of K's value.

These points leave open the question of how the metaphysical background in any given case is to be decided. Slater's preference is to leave the matter to the special sciences making use of the kind in question. For, as he puts it, 'different disciplines may tolerate different degrees of flexibility in the clustering required by their respective kinds' (Slater [2015], p. 403). Such differences are likely to be founded upon different presuppositions about what counterfactual perturbations are relevant. Our point here has been that these are typically metaphysical presuppositions. However, nothing we have said here would tell against the thought that scientists themselves are still the best parties to make such decisions. Nevertheless—whoever is making them—the decisions in question are metaphysical ones. They are not dictated purely by empirical data, nor are they dictated purely by pragmatic considerations about the ways natural kinds are used.

Such decisions are often innocuous in appearance. Clarke ([2010], pp. 313–4) has nicely highlighted the role that metaphysical presuppositions about individuality and persistence play in the biological sciences:

> Measuring fitness requires that biologists count biological individuals [. . .] Some biologists are happy to count offspring, while others insist on counting grand-offspring so that unhealthy offspring aren't credited. Some people prefer to think about fitness as signifying potential, rather

> than actual offspring, but it will be assumed that the actual count will converge on the same figure when you average over enough organisms of the same type. Often the absolute figure will be converted to a relative one by comparing the number of offspring with the numbers produced by conspecifics. Either way, the person doing the counting needs to keep track of the rate at which organisms reproduce to leave individuals of the same type as themselves in the next generation.

This is the kind of metaphysical decision that we have in mind. Some of the decision is guided by experience of the individuals in question, and some of it is guided by views about the property of fitness. But some of it is also guided by metaphysical assumptions about what an individual is in the first place. Both the case of the lichen symbiont as well as that of the Portuguese man o' war illustrate how the evaluation of a purported kind can hinge on just such assumptions.

Some recent critics have offered sharper criticisms of Slater's account, while others have offered broader criticisms of this family of accounts of natural kinds. To clarify the position we have developed here, it will be helpful to consider some of these alternatives.

Martinez ([2017]) suggests Slater's approach to natural kinds is incomplete because it rejects discussion of the underlying mechanisms of their stability. He suggests that the grounds of stable property clusters are, if not an ineliminable part of what explains the epistemic success of scientific practices, at least an important one. He argues, drawing upon examples from cancer research, that scientific research relies not only on stability of the clusters identified by researchers, but also on knowledge of relevant (and different) causal genetic mechanisms found in patients in order to assess their best treatment options. As such, he suggests an epistemology-only account requires recognition of grounding claims in order to evaluate cliquish stability. We agree with much of Martinez's criticism. However, in our view, the lesson learnt from examining cases of epistemology-only accounts (especially Slater's and those Martinez examines) is not that we should add consideration of metaphysical grounds to our epistemology of natural kinds in order to properly represent their use in scientific practice. In contrast, our analysis of what is problematic about neglecting the metaphysics of natural kinds (for example, omitting discussion of the grounds of the property clustering) is that the epistemic value of a natural kind is contingent upon the metaphysical presuppositions that lie behind the use of that kind in practice. Yet this is compatible with the possibility that there just is no systematic or interesting contribution a philosophical account of the causes or grounds of natural kinds might make to our understanding of their epistemic significance.

More broadly, Lemeire ([2018]) argues that any 'epistemology-only' account of natural kinds will inevitably fail, either (i) because it cannot

distinguish natural from conventional kinds, or (ii) because it cannot explain why the criteria they propose make for natural kinds. It is easy to envision how a proponent of Slater's view might reply: (i) we distinguish natural from conventional kinds in a given discipline by asking the scientists who work in that discipline which are which; and (ii) this criterion is likely to be accurate because the scientists who work in a given discipline are best equipped to make that call. Call this the 'scientist-as-metaphysician' approach. This approach would satisfy both of Lemeire's desiderata. Our project here, however, has been to highlight a problem that remains even if one relies on such an approach. Indeed, we suspect that the problem we have outlined here lies at the root of the difficulties Lemeire raises for epistemology-only accounts. A common feature of such accounts is that they leave the metaphysics to the scientists. Thus, proponents of such accounts are all committed to the scientist-as-metaphysician approach. The fact that the problem we have raised remains even for that approach is a sign that it is the real difficulty lying beneath Lemeire's criticism: appealing to the scientist-as-metaphysician leaves us with the task of identifying when and why the metaphysical presuppositions the scientist makes are correct or incorrect.

As some of our earlier remarks should indicate, we are not opposed to the scientist-as-metaphysician approach (see for example Kendig [2016]; Kendig and Eckdahl [2017]). It has been recently (and capably) defended by authors such as Magnus ([2018]). Magnus ([2018], p. 1436) proposes a strongly naturalist position that relies on the skills of scientists to accurately pick out natural kinds: 'When scientists introduce a category, they are responding to the world—rather than making an arbitrary decision about how to use words—just insofar as the category is a natural kind'. However, the cases we have discussed here highlight the fallibility of scientists with respect to such metaphysical theorizing. Our examples show that our confidence here requires that we take these categories to successfully pick out natural kinds. But whether they pick out natural kinds depends in turn on whether the metaphysical assumptions the scientists make are correct.

5 Conclusion

We agree with Slater that much progress can be made in understanding natural kinds by examining their epistemic value. Part of what makes talking about natural kindness useful is that it identifies an epistemically valuable feature shared by those clusters of properties we identify as real natural kinds. We can justify inferences based on natural kinds because of their relative stability, which is reflected in the property of natural kindness. The question we have raised is: how does natural kindness do this while remaining neutral with regard to the underlying mechanisms of stability? We suggested

that reliable identification of what qualifies as a mechanism of stability underpins its epistemic value. Our initial worry was that any strongly independent epistemic account may be vulnerable both to false positives and to false negatives. Such cases suggest that the assessment of cliquish stability is not something whose fallibility can be judged independently of the metaphysical presuppositions of the investigator. If stability is identified and attributed to a cluster on the basis of a preconceived notion of what is responsible for that cluster's cliquishness, the account would judge it to be a natural kind. The above examples of false positives and false negatives suggest this is a problem. Given the account of natural kinds Slater proposes, strong independence is not an option but weak independence is.

If Slater's aim is to avoid a metaphysical account of natural kinds, then it would seem he should prefer strong independence. However, if an epistemic account relies on weak independence and not strong independence, then the explanatory work that cliquish stability does in picking out the property of natural kindness is (at least in part) grounded in our knowledge that the entity has already been discovered to be a natural kind and so—by virtue of that knowledge—is the appropriate sort of entity that we can attribute natural kindness to. This of course is not only a problem Slater must face, but one that must be faced by all who seek a metaphysically neutral approach to explain the nature of natural kinds (for example, Khalidi [2013]; Magnus [2014], [2018]; Ereshefsky [2018]; Ereshefsky and Reydon [2015]). In order to secure a metaphysically neutral account of natural kinds that does not rely on kinds that have already been determined to be natural and stable, we must be able to identify mistakes that arise when the wrong conditions of stability for putative natural kinds are prescribed. Because the underlying metaphysical views that inform an investigator's decision to search for certain sorts of stabilizing mechanisms are often implicit, we also need to be able to identify mistakes in the attribution of natural kindness on the basis of tacitly held but not explicitly stated notions. Doing so provides a way to distinguish between inferences generated through the assumption of conditions of stability that are (or turn out to be) erroneous despite relying on good methodologies.

While we have not provided a solution to this problem here, our analysis suggests a potential approach to identifying such mistakes. This approach relies on assessing the epistemic status of the initial identification of cliquish stability. In discussing whether or not a cluster of properties can be understood to be the bearer of cliquish stability, we can ask questions about both the object of investigation, the putative natural kind grouping, as well as the subject who is grouping. For example, we might ask: is there a way to tell whether or not investigators lack the epistemic authority to make an attribution of cliquish stability even if the mechanisms they use are consistent with

the questions they ask and arise from sound methodological practices? Without knowing the implicit metaphysical commitments that the investigators use, we are unable to know the epistemic criteria for what sorts of things are potential candidates to anchor our inferences and who is a qualified 'anchorer'. Authors such as Slater who seek a purely epistemological account assume that we can determine what considerations are relevant to the identification of kinds without worrying about the metaphysics involved. Doing so assumes that the investigator's goals not only fully determine which mechanisms of stability are sought, but also that these are not open to criticism. Yet, as Haslanger ([2016], p. 131) points out, these are all open to investigation: 'The practice [of theorizing]—its ends and what it employs as means—is also open to critique'.

Acknowledgements

We warmly thank our colleagues in the MSU Metaphysics Reading Group (Matt McKeon, Emily Katz, David Godden, and Jason Mask) for open, challenging, and productive discussions. Impetus for writing this article began as a result of our initial (and perhaps less successful) attempts to respond to some of the more persistent and tenacious questions that arose in our exchanges.

Catherine Kendig
Department of Philosophy
Michigan State University
East Lansing, MI, USA
kendig@msu.edu

John Grey
Department of Philosophy
Michigan State University
East Lansing, MI, USA
greyjohn@msu.edu

References

Bardi, J. and Marques, A. C. [2007]: 'Taxonomic Redescription of the Portuguese Man-of-War, *Physalia physalis* (Cnidaria, Hydrozoa, Siphonophorae, Cystonectae) from Brazil', *Iheringia: Série Zoologia*, **97**, pp. 425–33.

Beschel, R. [1961]: 'Dating Rock Surfaces by Lichen Growth and Its Applications to Glaciology and Physiography (Lichenometry)', in G. O. Raasch (*ed.*), *Geology of the Arctic*, Toronto: University of Toronto Press, pp. 1044–62.

Boyd, R. N. [1991]: 'Realism, Anti-foundationalism, and the Enthusiasm for Natural Kinds', *Philosophical Studies*, **61**, pp. 127–48.

Boyd, R. [1999]: 'Homeostasis, Species, and Higher Taxa', in R. A. Wilson (*ed.*), *Species: New Interdisciplinary Essays*, Cambridge, MA: MIT Press, pp. 141–86.

Boyd, R. [2000]: 'Kinds as the "Workmanship of Men": Realism, Constructivism, and Natural Kinds', in J. Nida-Rümelin (*ed.*), *Rationalität, Realismus, Revision*, Berlin: De Gruyter, pp. 52–89.

Boyd, R. [2010]: 'Realism, Natural Kinds, and Philosophical Methods', in H. Beebee and N. Sabbarton-Leary (*eds*), *The Semantics and Metaphysics of Natural Kinds*, New York: Routledge, pp. 212–34.

Chagnon, P.-L., U'Ren, J. M., Miadlikowska, J., Lutzoni, F. and Elizabeth Arnold, A. [2016]: 'Interaction Type Influences Ecological Network Structure More Than Local Abiotic Conditions: Evidence from Endophytic and Endolichenic Fungi at a Continental Scale', *Oecologia*, **180**, pp. 181–91.

Clarke, E. [2010]: 'The Problem of Biological Individuality', *Biological Theory*, **5**, pp. 312–25.

Craver, C. F. [2009]: 'Mechanisms and Natural Kinds', *Philosophical Psychology*, **22**, pp. 575–94.

Ereshefsky, M. [2018]: 'Natural Kinds, Mind Independence, and Defeasibility', *Philosophy of Science*, **85**, pp. 845–56.

Ereshefsky, M. and Reydon, T. A. C. [2015]: 'Scientific Kinds', *Philosophical Studies*, **172**, pp. 969–86.

Green, T. G. A., Nash, T. H. and Lange, O. L. [2010]: 'Physiological Ecology of Carbon Dioxide Exchange', in T. H. Nash (*ed.*), *Lichen Biology*, Cambridge: Cambridge University Press, pp. 154–83.

Haslanger, S. [2016]: 'Theorizing with a Purpose: The Many Kinds of Sex', in C. Kendig (*ed.*), *Natural Kinds and Classification in Scientific Practice*, Abingdon: Routledge, pp. 129–44.

Henskens, F. L., Green, T. G. and Wilkins, A. [2012]: 'Cyanolichens Can Have Both Cyanobacteria and Green Algae in a Common Layer as Major Contributors to Photosynthesis', *Annals of Botany*, **110**, pp. 555–63.

Herrnstein, R. and Murray, C. [1994]: *The Bell Curve: Intelligence and Class Structure in American Life*, New York: Free Press.

Honeggar, R. [2010]: 'Morphogenesis', in T. H. Nash (*ed.*), *Lichen Biology*, Cambridge: Cambridge University Press, pp. 71–95.

Kappen, L., Schroeter, B., Green, T. G. A. and Seppelt, R. D. [1998]: 'Chlorophyll a Fluorescence and CO_2 Exchange on Umbilicaria Aprina under Extreme Light Stress in the Cold', *Oecologia*, **113**, pp. 325–31.

Kendig, C. [2016]: 'Homologizing as Kinding', in C. Kendig (*ed.*), *Natural Kinds and Classification in Scientific Practice*, Abingdon: Routledge, pp. 106–25.

Kendig, C. and Eckdahl, T. T. [2017]: 'Reengineering Metaphysics: Modularity, Parthood, and Evolvability in Metabolic Engineering', *Philosophy, Theory, and Practice in Biology*, **9**, pp. 1–21.

Khalidi, M. A. [2013]: *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*, Cambridge: Cambridge University Press.

Kirkpatrick, P. A. and Pugh, P. R.[1984]: *Siphonophores and Velellids*, London: Brill.

Kurlansky, M. [2002]: '*Physalia physalis:* Portugeese Man-of-War', *Animal Diversity Web*, available at <animaldiversity.org/accounts/Physalia_physalis/>.

Lange, M. [2009]: *Laws and Lawmakers*, Oxford: Oxford University Press.

Lawrey, J. and Diederich, P. [2003]: 'Lichenicolous Fungi: Interactions, Evolution, and Biodiversity', *The Bryologist*, **106**, pp. 80–120.

Lemeire, O. [2018]: 'No Purely Epistemic Theory Can Account for the Naturalness of Kinds', *Synthese*, available at <doi.org/10.1007/s11229-018-1806-8>.

Magnus, P. D. [2014]: 'NK≠HPC', *The Philosophical Quarterly*, **64**, pp. 471–7.

Magnus, P. D. [2018]: 'Taxonomy, Ontology, and Natural Kinds', *Synthese*, **195**, pp. 1427–39.

Martinez, E. J. [2017]: 'Stable Property Clusters and Their Grounds', *Philosophy of Science*, **84**, pp. 944–55.

Nash, T. [2010]: *Lichen Biology*, Cambridge: Cambridge University Press.

Pringle, A. [2017]: 'Establishing New Worlds: The Lichens of Petersham', in A. Tsing and H. A. Swanson (*eds*), *Arts of Living on a Damaged Planet*, Minneapolis, MN: University of Minnesota Press.

Rambold, G. and Triebel, D. [1992]: *The Inter-lecanoralean Associations*, Berlin: J. Cramer.

Siegel, E. [2017]: 'The Real Problem with Charles Murray and *The Bell Curve*', *Scientific American Online*, available at <blogs.scientificamerican.com/voices/the-real-problem-with-charles-murray-and-the-bell-curve/>.

Slater, M. H. [2015]: 'Natural Kindness', *British Journal for the Philosophy of Science*, **66**, pp. 375–411.

Spribille, T., Tuovinen, V., Resl, P., Vanderpool, D., Wolinski, H., Aime, C., Schneider, K., Stabentheiner, E., Toome-Heller, M., Thor, G., Mayrhofer, H., Johannesson, H. and McCutcheon, J. [2016]: 'Basidiomycete Yeasts in the Cortex of Ascomycete Macrolichens', *Science*, **353**, pp. 488–92.

U'Ren, J., Miadlikowska, J., Zimmerman, N., Lutzoni, F., Stajich, J. and Arnold, A. [2016]: 'Contributions of North American Endophytes to the Phylogeny, Ecology, and Taxonomy of Xylariaceae (Sordariomycetes, Ascomycota)', *Molecular Phylogenetics and Evolution*, **98**, pp. 210–32.