How to Incorporate Non-Epistemic Values into a Theory of Classification

Abstract

Non-epistemic values play important roles in classificatory practice, such that philosophical accounts of kinds and classification should be able to accommodate them. Available accounts fail to do so, however. Our aim is to fill this lacuna by showing how non-epistemic values feature in scientific classification, and how such values can be incorporated into a philosophical theory of classification and kinds. To achieve this aim, we present a novel account of kinds and classification (the *Grounded Functionality Account*), discuss examples from biological classification where non-epistemic values play a decisive role, and show how this account accommodates the role of non-epistemic values.

1. Introduction

In the past two to three decades the role of values in science has developed into a major topic of research in the philosophy of science (e.g., Longino, 1983; 1990; 1996; Rooney, 1992; 2017; Intemann, 2001; Kincaid, Dupré & Wiley, 2007; Douglas, 2009; 2016; Elliott, 2017; Elliott & McKaughan, 2014; Elliott & Steel, 2017). A common distinction in these discussions is between *epistemic* (or *constitutive*) values and *non-epistemic* (or *contextual*) values (Longino, 1983; 1990; 1996). The former category of values consists of values that promote the epistemic aims of science; examples include empirical adequacy, simplicity, inter-theoretic consistency, and testability. These are values that can be invoked in decisions on how best to achieve knowledge about the world – for example, decisions on theory choice or on the preferable interpretation of theories or data. The latter category consists of moral, social, cultural, aesthetic and other values. While the traditional view of science was that science is and should be free of non-epistemic values (while epistemic values play an important role in science), today a prominent view is that both epistemic and non-epistemic values are unavoidable factors in scientific practice. Some authors even argue that non-epistemic values are not only unavoidable in science, but also *should* play a role in good science (Longino, 1990; Anderson, 1995; Intemann, 2001; Douglas, 2007; 2017), sometimes even overruling epistemic values (Elliott & McKaughan, 2014). In addition, some authors doubt the feasibility of a clear distinction between epistemic and non-epistemic values, and argue that what might be considered an epistemic value in one context might be considered a non-epistemic value in another context (Rooney, 1992; 2017; Longino, 1996).

This deeper analysis of how values feature in the sciences has not permeated all of philosophy of science, though. One area that lags behind in this respect is the philosophical treatment of classification in the sciences. The aim of the present paper is to fill this lacuna by showing (1) how non-epistemic values play important roles in scientific classification, and (2) how such values can be incorporated into a philosophical theory of classification and kinds.

We proceed as follows. In Section 2, we discuss the role of both kinds of values in theories of classification and kinds, and give examples that show how non-epistemic values play important roles in classificatory practice. In Section 3, we briefly present a recently developed account of natural kinds, the *Grounded Functionality Account*, which accommodates the roles of (epistemic as well as non-epistemic) values in science. In Section 4, we show how this account handles the examples discussed in Section 2 and better captures the role of non-epistemic values in classificatory practices than other accounts of natural kinds. Section 5 concludes.

2. How values affect classification

2.1. Epistemic values and natural kinds

The role of epistemic values in the context of scientific classification is widely acknowledged. Indeed, epistemic values often are taken as criteria for deciding whether a given class should be conceived of as a natural kind. Consider for example Bird & Tobin's (2018) influential encyclopedia article on the topic. According to Bird & Tobin, some of the typical criteria that are used to attribute natural kind status to groups of entities are: that the member entities of a natural kind share one or more (natural) properties amongst each other; that natural kinds serve as the basis for inductive inferences; that natural kinds feature in one or several laws of nature; that natural kinds form a hierarchy of nested groups; and, that natural kinds should be categorically distinct from other kinds of kinds. Note that these criteria all involve epistemic values: they rest on valuations of epistemic aspects of science, such as that scientific classifications should serve to ground inferences, or should not encompass cross-cutting groups. These criteria are widely used in various combinations to explain what natural kinds *are* and what makes them focal groups of scientific studies.

Indeed, the prominence of epistemic criteria can be seen in many contemporary philosophical theories of natural kinds. Boyd's (1991; 1999) influential Homeostatic Property

Cluster account of kinds, for example, places emphasis on inferences supported by causal mechanisms and explicitly defines natural kinds as kinds that support inferential practices. Magnus' (2012) account of kinds, too, turns on inductive and explanatory success. Slater's (2013; 2015) Stable Property Clusters account conceives of natural kinds as patterns that can be observed in nature. And by way of final example, Khalidi's (2013; 2018) account places its focus on the causal structure of nature, conceiving of natural kinds "those categories that enable us to gain knowledge about reality" (2013: xi) and that are supported by nodes in nature's causal network. The epistemic values in these accounts focus on obtaining knowledge about the world, such as establishing bases for inference or accurate representation of certain aspects of nature (for instance, stable patterns or nodes in nature's causal network).

Philosophical accounts of classification and natural kinds thus center on classifications satisfying various epistemic virtues. But they neglect non-epistemic values: none of the accounts discussed above mention non-epistemic values, which we think is an important lacuna in the philosophical literature on classifications. Non-epistemic values frequently affect scientific research involving classifications and natural kinds. Moreover, they also affect the products of such research, i.e., the classifications themselves (as we will show in Section 4). Let us turn to a concrete example from biological research.

2.2. Non-epistemic values and natural kinds

It is widely agreed that non-epistemic values can come into play at various stages of research (Elliott, 2017; Reiss & Sprenger, 2017). For example, non-epistemic values can come into play when deciding which of the available data sets are relevant when testing a hypothesis, when deciding between competing theories that all explain a certain set of phenomena, and when deciding which of the many possible research projects to pursue in the first place (Kitcher, 2001). Here, we focus on the stage of theory acceptance or rejection, for which the role of non-epistemic values has been extensively addressed in the literature.[1] We will examine how non-epistemic values can affect the acceptance or rejection of a hypothesis or theory regarding how the entities in a particular domain should be classified and how they affect the way in which a classification is constructed and formulated.

The case that we examine concerns the acceptance or rejection of particular species concepts as the basis for the classification of organisms.[2] Species concepts can be thought of as theories about classification: a species concept tells us what biological species *are* and what binds organisms together into species, and as such determines how organisms are to be classified. The Phenetic Species Concept, for example, defines species on the basis of the overall similarity between their member organisms and as clusters of similar organisms between which clear differences exist (Sneath & Sokal, 1973). The Biological Species Concept specifies that species are "groups of interbreeding natural populations that are reproductively isolated from other such groups" (Mayr, 1996: 264). The Phylogenetic Species Concept in one of its various versions says that a species is "the smallest diagnosable cluster of individual organisms within which there is a parental pattern of ancestry and descent" (Cracraft 1983: 170). As a final example, the General Lineage Concept states that "a species

---

[1] Two well-known arguments that have been identified in this context are the *underdetermination argument* (Longino, 1990; Intemann, 2005) and the *argument from inductive risk* (Rudner, 1953; Douglas, 2009; 2017). According to the underdetermination argument, empirical evidence and epistemic values underdetermine which theory or hypothesis to accept, and non-epistemic values fill in the gap and provide scientists with a basis for theory choice. According to the argument from inductive risk, non-epistemic values are required to determine the evidential threshold for the acceptance or rejection of a theory or hypothesis. In the present paper, we will not be concerned with these two arguments.

[2] The examples are also briefly discussed in a different context by Conix (2019: 31-32).

corresponds with a lineage segment bounded by certain critical events […] species are segments of population-level lineages" (De Queiroz, 1999: 53).

While some species concepts are much more prominent than others, and some are not in use at all, they feature in the literature as competing theories about species as classificatory groups. First, they give different accounts of the metaphysics of species: species are similarity classes of organisms; species are groups (i.e., systems) of populations connected within a system by interbreeding and kept separate from other systems by reproductive isolation; species are groups of organisms that are part of a genealogical network and as a group are recognizable; species are segments of the Tree of Life bounded by critical events (i.e., speciation and extinction events). Second, they give different criteria for allocating organisms to species: trait similarity; breeding relations; relations of descent, and so on.

As competing theories about the nature of species and the classification of organisms, different species concepts in many cases yield different, incompatible groupings of organisms into species. In such cases, biologists are faced with a problem of theory choice: they must decide which species concept to use and which to reject as unsuitable for their research. As we will show in the following example, non-epistemic values play a crucial role in such decisions.

The Phylogenetic Species Concept is one of the most widely used species concepts in biology, yet in several contexts of investigation there are discussions regarding its applicability. In conservation biology, for example, a number of authors reject the Phylogenetic Species Concept because they think its adoption has detrimental consequences for conservation practice (Agapow et al., 2004; Isaac et al., 2004; Frankham et al., 2012; Zachos et al., 2013). Two main reasons are typically mentioned, both connected to the problem of taxonomic inflation. Studies have shown that classifying organisms using the Phylogenetic Species Concept often yields considerably more species than when the Biological Species Concept is used: on average one obtains 48% more species with associated decreases in these species' population sizes and ranges (Agapow et al., 2004). Extreme examples concern the number of birds-of-paradise in Australasia, in which the increase was from around 40 to around 90 species (Cracraft, 1992) and of endemic bird species in Mexico, in which the increase was from 101 to 249 species (Peterson & Navarro-Sigüenza, 1999).

One obvious reason why such inflation of species numbers is detrimental for conservation efforts is that there will simply be many more species to preserve (while resources available for conservation efforts are limited) than when counts are lower. The second reason is that the species identified using the Phylogenetic Species Concept tend to be more vulnerable. As mentioned, they tend to occur in smaller and less widely ranging populations than species identified using the Biological Species Concept. But it is not only the small number of species' members that makes them more vulnerable, but also that small populations have less genetic variation than larger ones and thus suffer from inbreeding depression (Frankham et al., 2012). Such species tend to be at greater risk of extinction than species identified using the Biological Species Concept (Zachos et al., 2013). Thus, application of the Phylogenetic Species Concept in conservation biology confounds conservation efforts. As Frankham et al. conclude, using the Biological Species Concept "will typically yield a classification appropriate to conservation concerns. Conversely, use of the diagnostic phylogenetic […] species concept […] will often lead to inappropriate classifications" (2012: 30).

Note that the reasons why authors in conservation contexts prefer the Biological Species Concept over the Phylogenetic Species Concept are not connected to the epistemic aims of science, but to practical, non-epistemic interests: given that we value species and have an interest in their conservation, the Biological Species Concept will constitute a better tool to achieve our aims than the Phylogenetic Species Concept. In different contexts, similar considerations may give rise to different results. In the context of epidemiology and public health research, for example, Attenborough (2015) argues that the Phylogenetic Species

Concept (PSC) does a better job than other species concepts at highlighting cryptic or incipient mosquito species. Identifying such cryptic groups of mosquitoes (which are extremely difficult to distinguish morphologically) is important for public health purposes, because some such groups can be the bearers of malaria parasites while others do not bear parasites and thus are of no concern. Attenborough concludes "that a fine-grained taxonomy, based on the PSC criterion of fixed inherited differences, and including recognition of cryptic and incipient species that are barely distinguishable or indistinguishable morphologically, is an important prerequisite of further fundamental biological research on these mosquito populations. Optimum practical intervention also depends upon it: in this case, not in a conservation context but to improve human health […]" (2015: 147). Here, too, non-epistemic values play a decisive role in choosing the preferred species concept. But they lead to a different conclusion than in conservation contexts.

Philosophically, what is going on in these debates on species concepts can be understood in terms of two different reasons that can underwrite the acceptance or rejection of theories: *truth* and *significance* (Anderson, 1995; Kitcher, 2001). Anderson pointed out that scientific investigations may pursue a multitude of aims, many of which are connected to practical human interests and non-epistemic values. Which questions are being asked – which are deemed significant – thus in part depends on practical interests and non-epistemic values, and so do the theories that are formulated to answer them. As she put it, "[m]any of the questions we ask science to answer come from the social context of science, not from its internal puzzle-generating activities. […] questions based on contextual interests require answers expressed in *terms that track those interests*." (Anderson, 1995: 53; emphasis added). In a similar vein, on Kitcher's notion of well-ordered science, science cannot be aimed at unveiling the truth about the world *simpliciter*, as there are far too many things that could be investigated, and there are too many truths that aren't particularly important to find out about. Rather, Kitcher argued, science should be aimed at finding out significant truths, where "what counts as significant science must be understood in the context of a particular group with particular practical interests and with a particular history" (Kitcher, 2001: 61). That is, on Kitcher's view scientific practice always involves decisions based on non-epistemic values (pertaining to what is significant to particular groups of people with particular interests) as well as epistemic values (pertaining to how those significant truths are indeed truths about the world).

Truth (or rather, empirical adequacy) *does* play a role in the examples discussed above: both critics and promoters of the Phylogenetic Species Concept (in the conservation case as well as the malaria case) think it captures real taxonomic units in the world, and in that way, that it is truthful. But because of this agreement, truth does not play a *decisive* role in the discussions. The decisive factor is the perceived significance of species identified using the Phylogenetic Species Concept in a particular context of research with concrete aims that are different from finding out "the true" classification of organisms into species. For researchers in the context of conservation, such species are too small and too vulnerable and thus drive up conservation costs. So, they are not significant units for conservation. For researchers aiming to combat malaria and promote human health species identified using the Phylogenetic Species Concept *are* significant for their purposes. The Phylogenetic Species Concept, then, is not being accepted or rejected on the basis of empirical accuracy or truthfulness, but according to whether the information it captures about the world is significant in a particular context of scientific practice. And it is moral and social values related to nature conservation and human health that determine significance in these contexts.

The examples discussed in this section show that non-epistemic values enter into classificatory practices in science, and that they can play decisive roles there. For the philosophy of classification and kinds this means that we need an account of classification and kinds that is able to accommodate both the epistemic and non-epistemic values that scientists employ when deciding on the theoretical basis for the classification of entities in a domain of study and when constructing classifications. We have noted in Section 1 that available accounts of natural kinds fail to meet this criterion. In the next section we present an account that does meet this criterion – the Grounded functionality Account of natural kinds (or, GFA) – before showing in Section 4 how the GFA handles the cases discussed above.

3. The Grounded Functionality Account of natural kinds

The GFA and its background considerations have been presented in detail elsewhere (Ereshefsky & Reydon, 2015; in press; Reydon, in press), so we will only provide a brief overview here.

The GFA is motivated primarily by the aim to adequately accommodate the various aims that scientists have when classifying the entities they study into kinds. This is partly in line with traditional views of natural kinds, according to which natural kinds are those kinds that stand at the focus of scientific investigations (Bird & Tobin, 2018). However, as was highlighted in Section 2.1 and elsewhere (Ereshefsky & Reydon, 2015, in press), while available philosophical accounts of natural kinds often focus on kinds in the sciences, they tend have too narrow a view of what the role of kinds is in the sciences. They tend to focus on specific epistemic aims that natural kinds and the classifications in which they feature are supposed to promote, such as capturing the causal structure of the world (Khalidi, 2013; 2018) or serving as a basis for inductive claims (Boyd, 1991; 1999; Magnus, 2012; Slater, 2013; 2015). But scientists posit natural kind classifications to achieve a wide variety of aims: these include the ones just mentioned but also aims like achieving a stable ordering system that can serve as a reference system for an area of work, achieving a classification that adequately represents evolutionary history, grouping entities in ways that describe their observable properties best, achieving kinds suitable to control natural phenomena, obtaining kinds that are useful for practical purposes such as managing and combating diseases, and so on (Ereshefsky & Reydon, 2015). Often, several aims will be intertwined in a particular research context, but different contexts (different areas of science, but also different academic communities at different times) will tend to focus on different sets of aims. Science is not *all* about making inferences, or about representing causes.

Note that we do not advocate for any specific set of legitimate classificatory aims in the sciences. Rather we are advocating for a thorough naturalism according to which a theory of natural kinds and classifications should be able to accommodate any classificatory aims that feature successfully in the sciences. By focusing on just one or a few epistemic aims (such as supporting inferences or representing the world's causal structure), available accounts of natural kinds are insufficiently naturalistic. Most importantly in the context of the present paper, they ignore non-epistemic aims that play important roles in the sciences.

The GFA is designed as an account of kinds and classification that captures the various epistemic as well as non-epistemic reasons scientists have for positing natural kind classifications. This makes the GFA a practice-oriented account: it is intended to capture the various actual classificatory practices found in the sciences and the aims featuring in them (whichever those might be), rather than provide us with an ideal model of kinds and classification that stands at a distance from scientific practice.

A central element that makes the GFA a naturalistic account is the notion of 'classificatory programs' (Ereshefsky, 2001; Ereshefsky & Reydon, 2015; in press). Classificatory programs are those parts of research programs in which classifications are constructed of the entities under study. They consist of sorting principles, motivating principles, and classifications. Sorting principles sort entities into kinds, and motivating principles are the epistemic and/or non-epistemic aims of research that motivate why the entities under study should be sorted in a particular way. The classifications produced by a classificatory program we hold (in line with Bird & Tobin, 2018) identify putative natural kinds. The background consideration is that if a classificatory program is successful in producing a classification that is useful for research activities, its success can be explained as successfully having identified groupings that represent relevant aspects of the world. Which aspects those are is left open in the GFA and differs among research contexts and classificatory programs. As an overarching account of kinds, the GFA says that successfully used scientific kinds represent aspects of the world, but does not prescribe which aspects those must be.

The GFA is a normative account of kinds and classifications in that it puts two constrains on classifications and the kinds that feature in them: the *Functionality Condition* and the *Grounding Condition*. According to the *Functionality Condition*, a classification should satisfy the specific aims it is posited for. In other words, a classification is constructed in the context of a classificatory program for a particular purpose or purposes (i.e., to perform certain functions in the context of a particular research setting), so if we are to accept it as a classification that groups entities into natural kinds, it should satisfy those particular functions. For instance, *if* a classification is posited to facilitate induction, then it should satisfy that aim; *if* it is posited to promote stability in classification, it should do that. The *Grounding Condition* requires that the way in which a classification satisfies its function should be grounded in the world and not *merely* our interests and our conceptions of what the world is like. The idea is that successful scientific classifications identify groupings that depend on relevant aspects of the world and not *merely* on the interests and actions of human beings (Ereshefsky & Reydon, in press). The Grounding Condition thus poses a restriction on natural kinds next to the Functionality Condition. The latter condition tells us which scientific groupings are potential natural kinds (those groupings that successfully serve the aims they were posited for). The Grounding Condition adds that from the set of groupings highlighted by the Functionality Condition, some groupings should be thought of as natural kinds so long as the relevant area of science gives us an account of how they are grounded in the world. In sum, the GFA tells us that those groupings that meet the Functionality Condition as well as the Grounding Condition should be considered to be natural kinds.

Consider the following brief example, the Biological Species Concept. Its aim is to posit groups of organisms that are evolutionary units (groups of organisms that evolve in tandem), within the research context that aims to explain how speciation occurs and how species come into being and remain in existence for extended periods of time as separate entities (Mallet, 2010; 2013: 682). The Biological Species Concept satisfies the Functionality Condition *if* it successfully sorts organisms into evolving systems of populations. It satisfies the Grounding Condition *if* the way the Biological Species Concept sorts organisms (in this case, according to successful interbreeding within a species and reproductive barriers among species), is indeed one way the world causes organisms to cluster together in evolutionary units.

We want to highlight two aspects of the GFA. First, satisfying the Functionality and Grounding Conditions is a *local* matter: the two conditions should be satisfied in relation to the specific aims of a classificatory project, and those differ among programs. In contrast to other philosophical accounts of classification, the GFA does not assume some overarching aim or set of aims for all classifications. This is in agreement with actual scientific practice. Many species concepts, for instance, were intended to be applicable only to a specific range of organisms, such as sexually reproducing organisms, or only birds, rather than throughout the whole of biodiversity. Accordingly, their success should be measured only with respect to their intended application. Second, the Functionality Condition and the Grounding Condition work together, as the aims of a classificatory program determine how a classification should be grounded in the world. For example, the Biological Species Concept aims to sort organisms into evolutionary units according to interbreeding. The Functionality Condition is satisfied if the Biological Species Concept sorts organisms into evolutionary units. The Grounding Condition is satisfied if interbreeding in the world does indeed result in evolving groups. Both conditions are in fact satisfied, so according to the GFA the Biological Species Concept identifies natural kinds.[3]

Let us now return to the examples discussed in Section 2, where we highlighted the role of non-epistemic values in decisions on the acceptance or rejection of the Phylogenetic Species Concept for the purposes of conservation and the promotion of human health. In the following section, we deepen our analysis of these examples and show how the GFA handles the acceptance or rejection of classifications according to non-epistemic values. By making comparisons with other accounts of kinds and classification, we show how the machinery of the GFA works.

4. The GFA in scientific practice

4.1. How the GFA handles decisions on species concepts

The two conditions contained in the GFA can serve as a tool for analysis of specific cases, such as the debates about the acceptance or rejection of specific species concepts for the purposes set in particular contexts of research. The question here would be how well competing species concepts meet the two conditions.

---

[3] One might wonder if this result runs afoul of the thesis that species are individuals (Hull, 1978). However, the GFA does not try to determine if a group of entities forms a spatiotemporally continuous individual or a class of spatially unrestricted entities. All the GFA is trying to capture is those units found in successful scientific classification, regardless of whether those units are best seen metaphysically as individuals or classes. The GFA attempts to remain neutral about such metaphysical issues, except that natural kinds should be grounded in the world. According to the GFA, natural kinds are simply those kinds that feature successfully in science and that success is explained in terms of the Functionality and Grounding Conditions.

Let us first examine the discussion in the context of conservation. Participants in the debate generally agree that the Phylogenetic Species Concept is favored over the Biological Species Concept as well as other species concepts *if* only diagnosability and wide applicability would be considered as criteria. The Phylogenetic Species Concept generally yields very fine-grained diagnosable groupings, because it uses unique traits (character states that have arisen in one branch of the Tree of Life) in combination with common descent to sort organisms into groupings (Reydon & Kunz, 2019).[4] The uniqueness of the traits involved guarantees diagnosability of the ensuing groups. In favor of the Phylogenetic Species Concept, it often meets the diagnosability requirement better than the Biological Species Concept, as groups defined by reproductive isolation are very difficult to diagnose in the wild or in the laboratory (Agapow et al., 2004: 163). In addition, authors note that the Phylogenetic Species Concept is applicable throughout the whole of biodiversity, which sets it apart from other species concepts (Agapow et al., 2004: 163). The Biological Species Concept, for example, is applicable only to sexually reproducing organisms. Thus, the Phylogenetic Species Concept performs better than competitor concepts on two important epistemic values – diagnosability and applicability.

Despite the Phylogenetic Species Concept satisfying these epistemic values, we saw in Section 2 that some conservation biologists reject the Phylogenetic Species Concept in conservation contexts for non-epistemic reasons. As we have shown in Section 2, using the Phylogenetic Species Concept tends to yield too many species, which we simply cannot all conserve because of limited resources, and it yields species that are too small and thus become threatened too easily. What is important to note is that in the debate we find a trade-

[4] The Phylogenetic Species Concept exists in several versions. The version quoted in Section 2 is sometimes called the "diagnosable version" (Mayden, 1997: 405), but the point we make holds for all versions.

off between diagnosability and applicability on the one hand, and practicability on the other. While both sides are clearly important for conservation purposes, in this trade-off between epistemic and non-epistemic values the latter are taken by participants in the debate to override the former.

For philosophical accounts of kinds and classification this case means that non-epistemic values should not only be allowed to play a role in classifications (for example, as factors operating in the background), but the account should also explicitly make room for the possibility that non-epistemic values override epistemic values. This implies that an adequate account of kinds and classification must treat epistemic and non-epistemic values on a par. Given their strong (or even exclusive) focus on epistemic values, available accounts of natural kinds fail to meet this condition. In what follows, we want to show that the GFA does meet this condition, comparing the GFA's performance to that of some other accounts.

The GFA reconstructs the debate on species concepts in conservation contexts as follows. As it meets the applicable epistemic requirements, the Phylogenetic Species Concept seems to meet the GFA's Functionality Condition, such that the groups based on it apparently could be highlighted as *putative* natural kinds (where we would have to invoke the Grounding Condition to decide whether they actually can be given natural kind status). However, adding non-epistemic values into the picture, we see why the Functionality Condition is in fact not met by the Phylogenetic Species Concept. While diagnosability furthers the aims of conservation to *some* extent (as a minimal level of diagnosability is required for conservation efforts to be effective in the first place), it yields too many diagnosable groups. In terms of the GFA, the opponents of the Phylogenetic Species Concept (Frankham et al., 2012; Zachos et al., 2013) argue that the Functionality Condition fails to be met by the Phylogenetic Species Concept, as it does not yield groupings of organisms that can stand at the focus of successful conservation efforts.

Had we only considered epistemic aims in our analysis of the debate, the conclusion would have had to be that the Phylogenetic Species Concept meets the Functionality Condition well and that Frankham and co-authors, and Zachos and co-authors were wrong in their rejection of the Phylogenetic Species Concept. Considering epistemic and non-epistemic aims on a par allows us to give an appropriate reconstruction of the debate and fundamentally changes the conclusion regarding the rejection of the Phylogenetic Species Concept. Note that in this analysis, the Functionality Condition has normative force, as it forces us to ask what the *actual* aims of a particular research context are. In the case considered here, the main aim is the conservation of biodiversity, while the epistemic aims of diagnosability and applicability are subordinate to it. This allows us to see that Frankham and co-authors, and Zachos and co-authors rightly reject the Phylogenetic Species Concept (*if* their empirical claims are correct), because they highlight the actual aims of the program in which the classification is supposed to be used. As the Functionality Condition is not met by the Phylogenetic Species Concept in this case, the Grounding Condition doesn't even come into play. After all, the GFA tells us to first test whether the Functionality Condition is met, and second whether the Grounding Condition is met.[5]

Let us now turn to the second example discussed in Section 2 – the case of the Phylogenetic Species Concept in the context of research on human health. As this is a different context of research and aims and values are context dependent, we might reach a

[5] Let us reiterate what the two conditions do. In the first step of the analysis, the Functionality Condition is used to identify successful kinds in science that are candidates for being thought of as natural kinds. In the second step the group of candidates is narrowed down, as the Grounding Condition is used to identify those kinds of which the successful use can be explained by their being grounded in relevant aspects of the world. Thus, the GFA can be used to identify natural kinds *and* to explain their successful use in scientific research.

different conclusion here. Note first that (as pointed out in Section 2) Attenborough (2015) argues that in the context of epidemiology when it comes to the prevention and eradication of malaria the capability of diagnosing cryptic species counts in favor of the Phylogenetic Species Concept. Using the Phylogenetic Species Concept, we can achieve a more fine-grained taxonomy of species in the genus *Anopheles*, which allows for more differentiated associations between the local presence of a particular type of mosquito and the connected risks for human health than on the basis of the traditionally recognized species. As Attenborough shows (2015: 144; Attenborough's Table 7.2), using the Phylogenetic Species Concept allows us to split two recognized species of *Anopheles* mosquitos into seven species that have different ranges of occurrence and different population densities, and that differ among each other with respect to whether their member organisms can carry malaria parasites and whether they prey on humans or on non-human animals.[6] Here, too, diagnosability is what counts. But because of the differences between the species that are recognized using the Phylogenetic Species Concept with respect to where they occur, whether they prey on humans, and whether they carry malaria parasites, the epistemic aim of diagnosibility in this case aligns with the non-epistemic aim of the promotion of human health in regions where malaria occurs.

In our analysis of this example, too, the Functionality Condition exerts normative force. It tells us to prioritize the actual aims of the research program under consideration (i.e., malaria epidemiology, where the aim is to promote human health). Following this guideline, we see that the Phylogenetic Species Concept fits the actual aims of this specific context better than other species concepts (assuming that the scientific claims in Attenborough's

[6] *A. punctulatus* (occurring in Pacific regions) becomes split into *A. punctulatus*, *A. farauti*, *A. hinesorum*, and *A. irenicus*. *A. gambiae* (occurring in Africa) becomes split into *A. gambiae*, *A. arabiensis*, and *A. quadriannulatus*.

paper are correct). The Grounding Condition adds to this by requiring that the groups identified by the Phylogenetic Species Concept represent aspects of the world that are relevant to the aims in malaria research. In this case, the aspects of the world are shared traits among mosquitoes, due to common descent, that pertain to whether or not they are able to carry malaria parasites and whether or not they prey on humans. The Phylogenetic Species Concept rests on common descent as well as unique traits, and as such highlights shared traits among species' members that explain why different groups of mosquitoes have different relevance when it comes to malaria research and prevention. The Grounding Condition thus is met in this case, which in turn explains why the Functionality Condition is met. The analysis thus allows us to see that Attenborough is right in endorsing the PSC (*if* his empirical claims are correct), because he highlights the actual aims of the program in which the classification is supposed to be used.

4.2 The GFA in comparison

In both examples we see that the GFA can make sense of the authors' argumentation. It can do this because it does not assume that all good scientific classifications further the same epistemic aim or aims, but instead examines classifications at a local level, taking the aims of a local classificatory program as the basis for analysis and treating epistemic and non-epistemic aims as being on a par. In this respect the GFA contrasts strongly with other prominent accounts. To clarify the contrast, let us briefly look at Khalidi's (2013; 2018) "causal nodes" account and Slater's (2013; 2015) "Stable Property Cluster" account as two examples.

On Khalidi's account, legitimate natural kinds are groupings of entities that represent nodes in causal networks (i.e., in the causal network structure of the world). Khalidi (2013; 2018) holds that classification should represent such nodes. On Slater's account, legitimate natural kinds are groupings of entities that represent stable patterns that we find in the world –

that is, stably recurring patterns of similarity between entities. Because patterns can be stable to larger and lesser degrees, on Slater's account kinds can be more or less good. Slater (2013; 2015) expresses this with his notion of "natural kindness" – being a natural kind, that is, natural kindness, comes in degrees. Note that these are *global* criteria that all kinds in all contexts must meet, and that these criteria do not differentiate between kinds that meet them. On Khalidi's account any classification that zooms in onto nodes in the world's causal nexus is as good as any other. If a grouping represents a node in the world's causal nexus, it should be given natural kind status. In a similar way, on Slater's account, any stable pattern is as good as any other. If a grouping of entities represents a stable pattern of property co-occurrence, it should be given natural kind status (albeit that in Slater's account, this status can be attributed to a certain degree). What the two accounts lack are filters that would allow us to distinguish between more and less *significant* causal nodes and more and less *significant* stable patterns, respectively.

This entails that on both Khalidi's and Slater's accounts we would have to conclude that Frankham and co-authors, and Zachos and co-authors are wrong when they argue against the use of the Phylogenetic Species Concept in conservation contexts. After all, the Phylogenetic Species Concept works perfectly well when it comes to identifying causal nodes (here inheritance from a common ancestor can be interpreted as a causal node) or stable patterns (shared traits that uniquely define a group, even if not all members actually exhibit the traits). Also, we would have to conclude that while Attenborough is right to prefer the Phylogenetic Species Concept in the context of malaria epidemiology, he is right for the wrong reasons: he should have emphasized diagnosability of groups based on causally sustained (genetic, morphological and behavioral) similarities, or on the basis of the fact that the Phylogenetic Species Concept is more powerful than competing concepts when it comes to identifying stable patterns in the world (as it recognizes more finely grained patterns). On Khalidi's and Slater's accounts of classification, Attenborough should not have preferred the Phylogenetic

Species Concept on the basis of whether it is able to individuate groups of mosquitos that are important for epidemiological research and the promotion of human health.

The crucial problem with both Khalidi's and Slater's accounts, we suggest, is that they miss the aims that in actual research contexts are set by the non-epistemic values that are endorsed by the community of researchers in these cases. As we have seen, these non-epistemic aims do not distinguish between different aspects of the causal structure of the world or between more and less relevant stable patterns – the aims highlighted in Khalidi's and Slater's account, respectively.[7] But in the two examples that we discussed in the present paper, the non-epistemic aims are the ones that matter the most and are taken by the researchers involved to override epistemic values. One important thing that we have seen by looking at two examples of actual scientific practice is that the GFA's Functionality Condition is not met overarchingly by one epistemic aim in all scientific contexts. Whether it is met depends on the research context and the goals set in that context. The assessment whether a particular classificatory theory (in the cases considered, a particular species concept) yields natural kinds has to be carried out locally, as kinds are strongly context dependent – which is an aspect of kinds that other accounts miss.

To be sure, the context dependency of classifications and the kinds that feature in them is acknowledged in some of the available accounts. Boyd's (1991; 1999) Homeostatic Property Cluster account, for example, explicitly conceives of kinds as relative to disciplinary

[7] In later work, Slater (2017) does introduce a normative aspect that allows us to make such distinctions. Slater suggests that "classificatory practices are laden with normative commitments of a distinctive kind" (2017: 1). The norms he includes in his account, however, all are practical norms (such as the norm that "lonely categories" – categories including only one element – are to be avoided). Such norms are not non-epistemic norms of the stripe described in this paper.

matrices (Boyd, 1999: 148). That is, for Boyd natural kinds are groupings that accommodate the inferential practices within a particular disciplinary matrix to the causal-mechanical structure of the world. Boyd does not think of disciplinary matrices as corresponding to scientific disciplines as these are commonly understood, but explains that a disciplinary matrix is "a family of inductive and inferential practices united by common conceptual resources, whether or not these correspond to academic or practical disciplines otherwise understood" (Boyd, 1999: 148). This, however, is not a thoroughly local context dependency, as the disciplinary matrices that Boyd refers to are typically located at comparatively high levels of organization, ranging over one or multiple disciplines. The GFA is much more fine-grained in this respect and aims to be *thoroughly local* (in line with Reydon, 2016; see also Conix, 2019: 33): it evaluates the status of kinds and classification according to the aims of local research contexts, which may for example be a specific school of thought or a concrete practical project within a small subdiscipline. Furthermore, Boyd's account exclusively takes induction and inference as the aims for which kinds and classifications are constructed, whereas the GFA is completely open with respect to the aims of a classificatory program.

Concluding this section, we have attempted to show how the GFA allows us to accommodate important aspects of scientific practice that available accounts of kinds and classification do not accommodate, and thus allows us to reconstruct the decisions that scientists make (in the examples discussed, decisions on theory choice) better than other accounts. These aspects are: that non-epistemic values play a role in classifications; that non-epistemic values must be treated on a par with epistemic values; and that sometimes non-epistemic values even override epistemic values, because the ultimate aims of some research contexts are non-epistemic.

5. Conclusion

The role of non-epistemic values in scientific classification has been a neglected topic in the philosophy of science. In the present paper, we have tried to show why this is a lacuna in philosophical research – because non-epistemic values sometimes play a decisive role in how scientists construct classifications – and presented an account of kinds and classification that can fill this lacuna.

Our account is thoroughly naturalistic and practice-oriented. It treats epistemic and non-epistemic values on a par, because in some cases researchers will take epistemic values to overrule non-epistemic values, whereas in other cases this will be the opposite. Thus, an adequate account of kinds and classifications should not generally prioritize one type of values over the other. As available accounts are heavily focused on epistemic values and either ignore non-epistemic values or treat them as less important background factors, we believe these accounts are not adequate to actual scientific practice. Our account, the GFA, aims to change the focus of philosophical thinking about scientific kinds and classification in this respect.

References

Agapow, P.-M., Bininda-Emonds, O.R.P., Crandall, K.A., et al. (2004): 'The impact of species concept on biodiversity studies', *Quarterly Review of Biology* 79: 161-179.

Anderson, E. (1995): 'Knowledge, human interests, and objectivity in feminist epistemology', *Philosophical Topics* 23: 27-58.

Attenborough, R. (2015): 'What are species and why does it matter? Anopheline taxonomy and the transmission of malaria', in: Behie, A.M. & Oxenham, M.F. (Eds): *Taxonomic Tapestries: The Threads of Evolutionary, Behavioural and Conservation Research*, Canberra: ANU Press, pp. 129-151.

Bird, A. & Tobin, E. (2018): ‘Natural kinds’, in: Zalta, E.N. (Ed.): *The Stanford Encyclopedia of Philosophy (Spring 2018 Edition)*, <https://plato.stanford.edu/archives/spr2018/entries/natural-kinds/>.

Boyd, R.N. (1991): ‘Realism, anti-foundationalism and the enthusiasm for natural kinds’ *Philosophical Studies* 61: 127-148.

——— (1999): ‘Homeostasis, species, and higher taxa’, in: Wilson, R.A. (Ed.): *Species: New Interdisciplinary Essays*, Cambridge, MA: MIT Press, pp. 141-185.

Conix, S. (2019): ‘Radical pluralism, classificatory norms and the legitimacy of species classifications’, *Studies in History and Philosophy of Biological and Biomedical Science* 73: 27-34.

Cracraft, J. (1983): ‘Species concepts and speciation analysis’, in: Johnston, R.F. (Ed.): *Current Ornithology, Volume 1*, New York and London: Plenum Press, pp. 159-187.

——— (1992): ‘The species of the birds-of-paradise (Paradisaeidae): Applying the phylogenetic species concept to a complex pattern of diversification’, *Cladistics* 8: 1-43.

De Queiroz, K. (1999): ‘The general lineage concept of species and the defining properties of the species category’, in: Wilson, R.A. (Ed.): *Species: New Interdisciplinary Essays*. Cambridge, MA: MIT Press, pp. 49-89.

Douglas, H.E. (2007): ‘’Rejecting the ideal of value-free science’, in: Kincaid, H., Dupré, J. & Wylie, A. (Eds): *Value-Free Science? Ideals and Illusions*, New York: Oxford University Press, pp. 120-139.

——— (2009): *Science, Policy, and the Value-Free Ideal*, Pittsburgh: University of Pittsburgh Press.

——— (2016): ‘Values in science’, in: Humphreys, P. (Ed.): *The Oxford Handbook of Philosophy of Science*, New York: Oxford University Press, pp. 609-630.

——— (2017): ‘Why inductive risk requires values in science’, in: Elliott, K.C. & Steel, D. (Eds): *Current Controversies in Values and Science*, New York & London: Routledge, pp. 81-93.

Elliott, K.C. (2017): *A Tapestry of Values: An Introduction to Values in Science*, New York: Oxford University Press.

Elliott, K.C. & McKaughan, D.J. (2014): ‘Nonepistemic values and the multiple goals of science’, *Philosophy of Science* 81: 1-21.

Elliott, K.C. & Steel, D. (Eds) (2017): *Current Controversies in Values and Science*, New York & London: Routledge.

Ereshefsky, M. (2001): *The Poverty of the Linnaean Hierarchy: A Philosophical Study of Biological Taxonomy*, Cambridge: Cambridge University Press.

Ereshefsky, M. & Reydon, T.A.C. (2015): ‘Scientific kinds’, *Philosophical Studies* 172: 969-986.

——— (in press): ‘The Grounded Functionality Account of natural kinds’, in: Bausman, W., Baxter, J., Lean, O., Love, A. & Waters, C.K. (Eds): *From Biological Practice to Scientific Metaphysics*, Minneapolis: University of Minnesota Press, in press.

Frankham, R., Ballou, J.D., Dudash, M.R., et al. (2012): ‘Implications of different species concepts for conserving biodiversity’, *Biological Conservation* 153: 25-31.

Hull, D. (1978): ‘A matter of individuality’, *Philosophy of Science* 45: 335-360.

Intemann, K. (2001): ‘Science and values: Are value judgments always irrelevant to the justification of scientific claims?’, *Philosophy of Science* 68: S506-S518.

——— (2005): ‘Feminism, underdetermination, and values in science’, *Philosophy of Science* 72: 1001-1012.

Isaac, N.J.B., Mallet, J. & Mace, G.M. (2004): ‘Taxonomic inflation: Its influence on macroecology and conservation’, *Trends in ecology and Evolution* 19: 464-469.

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

——— (2018): 'Natural kinds as nodes in causal networks', *Synthese* 195:1379-1396.

Kincaid, H., Dupré, J. & Wylie, A. (Eds) (2007): *Value-Free Science? Ideals and Illusions*, New York: Oxford University Press.

Kitcher, P. (2001): *Science, Truth, and Democracy*, New York: Oxford University Press.

Longino, H.E. (1983): 'Beyond "bad science": Skeptical reflections on the value-freedom of scientific inquiry', *Science Technology & Human Values* 8: 7-17.

——— (1990): *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*, Princeton: Princeton University Press.

——— (1996): 'Cognitive and non-cognitive values in science: Rethinking the Dichotomy', in: Hankinson Nelson, L. & Nelson, J. (Eds): *Feminism, Science, and the Philosophy of Science,* Dordrecht: Kluwer, pp. 39-58.

Magnus, P.D. (2012): *Scientific Enquiry and Natural Kinds: From Planets to Mallards*, Houndmills, Basingstoke: Palgrave Macmillan.

Mallet, J. (2010): 'Group selection and the development of the biological species concept', *Philosophical Transactions of the Royal Society B* 365: 1853-1863.

——— (2013): 'Species, concepts of', in: Levin, S. (Ed.): *Encyclopedia of Biodiversity (Second Edition), Volume 6*, Amsterdam: Academic Press, pp. 679-691.

Mayden, R.L. (1997): 'A hierarchy of species concepts: The denouement in the saga of the species problem', in: Claridge, M.F., Dawah, H.A. & Wilson, M.R. (Eds): *Species: The Units of Biodiversity*, London: Chapman & Hall, pp. 381-424.

Mayr, E. (1996): 'What is a species, and what is not?', *Philosophy of Science* 63: 262-277.

Peterson, A.T. & Navarro-Sigüenza, A.G. (1999): 'Alternate species concepts as bases for determining priority conservation areas', *Conservation Biology* 13: 427-431.

Reiss, J. & Sprenger, J. (2017): 'Scientific objectivity', in: Zalta, E.N. (Ed.): *The Stanford Encyclopedia of Philosophy (Winter 2017 Edition)*, <https://plato.stanford.edu/archives/win2017/entries/scientific-objectivity/>.

Reydon, T.A.C. (2016): 'From a zooming-in model to a co-creation model: Towards a more dynamic account of classification and kinds', in: Kendig, C.E. (Ed.): *Natural Kinds and Classification in Scientific Practice*, London & New York: Routledge, pp. 59-73.

——— (in press): 'Functional kinds and the metaphysics of functional levels: In what sense are functionally defined kinds and levels non-arbitrary?', in: Brooks, D., DiFrisco, J. & Wimsatt, W.C. (Eds): *Levels of Organization in the Biological Sciences*, Cambridge (MA): MIT Press, in press.

Reydon, T.A.C. & Kunz, W. (2019): 'Species as natural entities, instrumental units and ranked taxa: New perspectives on the grouping and ranking problems', *Biological Journal of the Linnean Society* 126: 623-636.

Rooney, P. (1992): 'On values in science: Is the epistemic/non-epistemic distinction useful?', in: Hull, D.L., Forber, M. & Okruhlik, K. (Eds): *PSA 1992: Proceedings of the Biennial Meeting of the Philosophy of Science Association (Volume One: Contributed Papers)*, East Lansing, MI: Philosophy of Science Association, pp. 13-22.

——— (2017): 'The borderlands between epistemic and non-epistemic values', in: Elliott, K.C. & Steel, D. (Eds): *Current Controversies in Values and Science*, New York & London: Routledge, pp. 31-46.

Rudner, R. (1953): 'The scientist qua scientist makes value judgments', *Philosophy of Science* 20: 1-6.

Slater, M.H. (2013): *Are Species Real? An Essay on the Metaphysics of Species*, Houndmills, Basingstoke: Palgrave Macmillan.

——— (2015): 'Natural kindness', *British Journal for the Philosophy of Science* 66: 375-411.

——— (2017): ‘Pluto and the platypus: An odd ball and an odd duck – On classificatory norms’, *Studies in History and Philosophy of Science* 61: 1-10.

Sneath, P.H.A. & Sokal, R.R. (1973): *Numerical Taxonomy*, San Francisco: W.H. Freeman and Company.

Zachos, F.E., Apollonio, M., Bärmann, E.V., et al., (2013): ‘Species inflation and taxonomic artefacts—A critical comment on recent trends in mammalian classification’, *Mammalian Biology* 78: 1-6.