



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

Thomas A. C. Reydon¹ · Marc Ereshefsky²

Received: 10 August 2020 / Accepted: 15 December 2021 / Published online: 8 January 2022
© The Author(s) 2022

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 they can be incorporated into a philosophical theory of classification and kinds. To achieve this, we present a novel account of kinds and classification (the *Grounded Functionality Account*), discuss examples from biological classification where non-epistemic values play decisive roles, and show how this account accommodates the role of non-epistemic values.

Keywords Natural kinds · Values · Non-epistemic values · Grounded functionality account of natural kinds · GFA

1 Introduction

In the past few decades the role of values in science has developed into a major topic in philosophy of science (e.g., Longino, 1983, 1990, 1996; Rooney, 1992, 2017; Intemann, 2001; Kincaid et al., 2007; Douglas, 2009, 2016; Elliott, 2017; Elliott & McKaughan, 2014; Elliott & Steel, 2017). A common distinction in the literature is between *epistemic* (or *constitutive*) and *non-epistemic* (or *contextual*) values (Longino, 1983, 1990, 1996). The former category consists of values that promote the epistemic aims of science; examples include empirical adequacy, simplicity,

✉ Thomas A. C. Reydon
reydon@ww.uni-hannover.de

Marc Ereshefsky
ereshefs@ucalgary.ca

¹ Institute of Philosophy & Centre for Ethics and Law in the Life Sciences (CELLS), Leibniz University Hannover, Im Moore 21, 30167 Hannover, Germany

² Department of Philosophy, University of Calgary, 2500 University Drive. N.W, Calgary, AB T2N 1N4, Canada

inter-theoretic consistency, and testability. These values can be invoked in decisions on, for instance, theory choice, the preferable interpretation of theories or data, etc. The latter category consists of moral, social, political, cultural, aesthetic and other values.

While on the traditional view science should be free of non-epistemic values (while recognizing epistemic values as important 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, 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 rich analysis of values in science has not permeated all of philosophy of science, however. One area lagging behind in this respect is the philosophical treatment of scientific classification and kinds. Some authors, such as Griffiths (1997: 200) and Khalidi (2013: 157ff.), have argued that non-epistemic values should *not* affect how scientists construct classifications into natural kinds.¹ Both think that natural kinds are posited to do epistemic work and that when values affect how we determine kinds, this detracts from the epistemic work kinds should do. In contrast, we believe that non-epistemic values are *as* important as epistemic values, and that an account of kinds and classification is needed that explicitly accommodates how non-epistemic values feature in classificatory practices.

This is not to say that the role of non-epistemic values in classification has not been discussed. Intemann (2015: 224), for example, considered the notion of ‘biodiversity’, arguing that how one measures biodiversity depends on what one values preserving. Considering an example from economics, Anderson (1995: 45-46) pointed out that for many years women were, on the basis of non-epistemic values, considered not to belong in the categories ‘employed’ and ‘unemployed’. If you were a woman and were “keeping house”, you were excluded from employment surveys – who was counted as employed or unemployed was dictated by an androcentric norm. Dupré (1993), Magnus (2012), Ludwig (2014, 2016), Brigandt (2015), and Conix (2018) all provided accounts of classification that are explicitly open to non-epistemic values playing a role in classificatory practices. And non-epistemic values in classification were also discussed by Hacking (1995, 2007) in his work on “looping effects”, and by authors debating psychiatric classification (Haslam, 2002a, 2002b; Zachar, 2014, 2015; Zachar & Potter, 2010).

Notwithstanding this attention for non-epistemic values in classification, available accounts of classification and kinds fail to explicitly incorporate the roles of non-epistemic values and remain strongly focused on epistemic values. The authors mentioned above, for instance, either discuss the role of non-epistemic values without

¹ In later work, however, Griffiths did make room for non-epistemic values in classification (e.g., Griffiths, 2004).

presenting a general account of how classification works, or (when they do have such an account, such as Dupré, Brigandt or Conix) do not recognize non-epistemic values as playing *decisive* roles in classificatory practices. The aim of the present paper is to fill this lacuna by showing (1) how non-epistemic values can play decisive roles in scientific classification (and override epistemic values), and (2) how such values can be incorporated into a recently developed philosophical theory of classification and kinds. In Section 2, we discuss how epistemic and non-epistemic values feature in theories of classification and kinds, and give examples that show how non-epistemic values play decisive roles in classificatory practice. In Section 3, we briefly present a recently developed account of natural kinds, the *Grounded Functionality Account* (GFA), which accommodates the roles of (epistemic as well as non-epistemic) values in science. In Section 4, we show how the GFA handles the examples discussed in Section 2 and better captures the role of non-epistemic values in classificatory practices than other available accounts. Section 5 concludes.

Before beginning, we want to clarify three things. First, we are sympathetic to the arguments that have been advanced against making a strict distinction between epistemic and non-epistemic values (Rooney, 1992, 2017; Longino, 1996). For example, Longino (1996) argues that the value of simplicity, which many hold is an epistemic value, is a non-epistemic value in the context of feminism, because diversity and ontological heterogeneity are taken as non-epistemic values that promote social justice. Nevertheless, because authors in the literature of science and values commonly use the terms ‘epistemic values’ and ‘non-epistemic values’ we will follow this terminology in this paper. Following the standard terminology in the literature on science and values makes it easier for us to show how so-called non-epistemic values can – and should – be incorporated into a theory of classification. But it is important to note that nothing in our paper hinges on this terminology and indeed our approach does not rely on any strong distinction between epistemic and non-epistemic values.

Second, we use the distinction between epistemic and non-epistemic values because this allows us to make our points most clearly. Available accounts of kinds and classification have neglected non-epistemic values and have not explicitly accommodated their roles in classificatory practices. In the cases we discuss, epistemic and non-epistemic values conflict, or epistemic values alone are insufficient to account for the choices scientists make in practice and non-epistemic values are decisive. This interplay of epistemic and non-epistemic values is an important but neglected aspect of classificatory practice, thus showing why a better account of classification and kinds is needed that accommodates the roles of all values in classificatory practice.

Third, we claim that the required account is already available, namely the GFA, and we aim to show how the GFA can accommodate the roles of epistemic as well as non-epistemic values in classificatory practice in ways other accounts cannot. While we discuss examples from biological science, we believe our claims generalize: non-epistemic values can play decisive roles in classificatory practices throughout the sciences, and as a general account the GFA is able to accommodate such cases. As the GFA has been developed in detail elsewhere (Ereshefsky & Reydon, 2015, 2022; Reydon, 2021) and space here is limited, we can only provide a brief introduction to the GFA and have to refer to other work for more details.

2 How values affect classification

2.1 Epistemic values and natural kinds

The role of epistemic values in scientific classification is widely acknowledged. Indeed, epistemic values often are taken as criteria for deciding whether a grouping should be conceived of as a natural kind. According to Bird & Tobin (2018), for example, the typical criteria used to attribute natural kind status to a group of entities include: its member entities share one or more (natural) properties amongst each other; it serves as the basis for inductive inferences; it features in laws of nature; it is part of a hierarchy of nested groups; and it is categorically distinct from other kinds. Note that these criteria all involve epistemic values: they rest on valuations of epistemic aspects of science – e.g., 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 values is seen in many contemporary theories of natural kinds. Boyd's (1991, 1999) influential Homeostatic Property Cluster (HPC) theory, for example, emphasizes inferences supported by causal mechanisms and explicitly defines natural kinds as those kinds that support inferential practices. Magnus' (2012) account, too, turns on inductive and explanatory success. Slater's (2013, 2015) Stable Property Cluster account conceives of natural kinds as patterns that can be observed in nature. And finally, Khalidi's (2013, 2018) account places its focus on the causal structure of nature, conceiving of natural kinds as "those categories that enable us to gain knowledge about reality" (2013: xi) and represent 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 accurately representing certain aspects of nature (stable patterns or causal nodes).

Philosophical accounts of classification and natural kinds thus center on various epistemic virtues classifications should satisfy. But they neglect non-epistemic values, such as social or moral values – none of the accounts discussed above mention such values. This is a lacuna, as non-epistemic values frequently affect classificatory practices and their products. Let us turn to a concrete example from biological research to show this.

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, they can feature in decisions concerning which of the available data sets are relevant when testing a hypothesis, decisions between competing theories that all explain a certain set of phenomena, and decisions concerning which of the many possible research projects to pursue in the first place (Kitcher, 2001). Here, we consider the stage of theory

acceptance or rejection, for which the role of non-epistemic values has been extensively addressed in the literature.² We examine how non-epistemic values can affect the choice between hypotheses (or theories) regarding the classification of entities in a particular domain, and how they affect the way in which classifications are constructed and formulated.

The two cases that we examine concern the acceptance of particular species concepts as the basis for the classification of organisms.³ 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 should be classified. The Biological Species Concept (BSC), for example, specifies that species are “groups of interbreeding natural populations that are reproductively isolated from other such groups” (Mayr, 1996: 264). The Phylogenetic Species Concept (PSC) 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). While some species concepts are much more prominent than others (and some not in use at all), they feature in the literature as competing theories about species as classificatory groups, giving different pictures of what species are: systems of populations connected by interbreeding and kept separate from other systems by reproductive isolation; or, groups of organisms that are part of a genealogical network and as a group are recognizable; etc. And they give different criteria for allocating organisms to species (Reydon & Kunz, 2019): breeding relations; relations of descent, etc. As competing theories about the nature of species and the classification of organisms, different species concepts often yield different, incompatible groupings of organisms into species. In such cases, biologists confront a choice problem: they must decide which species concept to use and which to consider less suitable or unsuitable for their research. As we show in what follows, non-epistemic values play a crucial role in such decisions.

The PSC is among the most widely used species concepts in biology, yet in several contexts of investigation there are discussions regarding its applicability. As a first example, consider conservation biology. Here, a number of authors are critical of the PSC because they think its adoption entails clear disadvantages for conservation practice (Agapow et al., 2004; Isaac et al., 2004; Frankham et al., 2012; Zachos et al., 2013). Frankham et al. argued that “the diagnostic phylogenetic species concept is unsuitable for use in conservation contexts” (2012: 25) and “[u]se of the

² Two well-known arguments 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 to 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. We will not be concerned with these two arguments here.

³ Ludwig (2014, 2016) and Conix (2017, 2018, 2019) also discuss these cases, but do not use them to support a general account of natural kinds. Ludwig and Conix use these examples to show how values can come into play in classificatory contexts in the sciences. We use their examples to provide support for our general account of natural kinds.

differential fitness or biological species concepts will typically yield a classification appropriate to conservation concerns. Conversely, use of the diagnostic phylogenetic or the taxonomic species concepts will often lead to inappropriate classifications.” (2012: 30). In a similar vein, Zachos et al. argued that “[t]he PSC [...] has serious shortcomings that make its application inappropriate on theoretical and practical grounds” (2013: 2). Other authors do not disqualify the PSC for use in conservation contexts, but do recommend caution in relation to its use. Isaac et al., for instance, emphasized that using the PSC entails effects that “will make global targets, such as the recent pledge in the Convention on Biological Diversity [...] to slow the loss of biodiversity by 2010, very hard to meet, as they confound our attempts even to measure the loss of species reliably” (2004: 466). Agapow et al. (2004) highlighted advantages as well as problems for all species concepts – they explicitly do not reject the PSC for use in conservation contexts but still are quite critical of its effects on conservation practice.⁴

Several disadvantages with the PSC are typically mentioned. Studies have shown that classifying organisms using the PSC often yields considerably more species than when the BSC is used: on average one obtains 48% more species with associated decreases in population sizes and ranges (Agapow et al., 2004; Isaac et al., 2004). Extreme examples concern the number of birds-of-paradise in Australasia, in which species numbers increased from around 40 to around 90 (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 taxonomic inflation is detrimental for conservation efforts is that there will simply be many more species to preserve (with limited resources available for conservation efforts) than with lower counts. But there will not only be more groups to conserve, they will also on average be more threatened. The species identified using the PSC tend to be more vulnerable to extinction than species identified using the BSC, as they tend to occur in smaller and less widely ranging populations (small populations do not only have smaller numbers of organisms but also less genetic variation than larger ones, making them prone to inbreeding depression – Frankham et al., 2012). Because such species tend to be at greater risk of extinction than species identified using the BSC (Zachos et al., 2013), application of the PSC confounds conservation efforts. As Frankham et al. conclude, using the BSC “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). Adding to this there is a social factor, namely that taxonomic inflation may lead to a demotivation of the general public to support conservation efforts. As Agapow et al. emphasized, “[p]ublic concern has a great influence over the allocation of conservation resources [...]

⁴ As they wrote, “recovering all species listed currently would cost around \$4.6 billion. With widespread adoption of the PSC, this already formidable amount could increase to \$7.6 billion, or the entire annual budget for the administering agency (U.S. Fisheries and Wildlife Services) for the next 120 years. [...] It has been said that the economic cost of current conservation strategies are already unacceptable [...]. The impact of the PSC serves only to reinforce this point” (Agapow et al., 2004: 169-170).

Changing species identities and a flood of threatened species may create confusion and apathy in the public's mind. Furthermore, consistent application of the PSC may make this task [education and mobilizing public opinion for conservation] more difficult by identifying morphologically disparate entities as members of the same species and superficially similar entities as members of different species" (2004: 170).

Note that the abovementioned reasons why authors in some conservation contexts prefer the BSC over the PSC are not connected to *epistemic* aims, but to non-epistemic interests: *given that* we value species and have an interest in their conservation, the BSC will constitute a better tool to achieve our aims than the PSC.⁵ In different contexts, similar considerations may give rise to different results. As a second example, consider the context of epidemiology and public health research. Attenborough (2015) argues that here the PSC performs better than other species concepts at highlighting cryptic and incipient mosquito species. Identifying such cryptic groups (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, these debates 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 various aims, many being connected to practical interests and non-epistemic values. Which questions are 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 writes, "[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). Similarly, 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 discover. Rather, Kitcher argued, science should be aimed at finding out *significant*

⁵ In response to a reviewer's comment, we should acknowledge that many conservation efforts are not focused on species, but on evolutionarily significant units that are considered worthy of conservation independently of the question whether or not such a unit has species status. In such contexts, the number of groups worthy of conservation clearly does not depend on the species concept that is being used. But our example is only intended to show that in *some* conservation contexts (i.e., those in which species status matters) non-epistemic values may override epistemic values. Other conservation contexts may well be different.

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, 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 supported empirically).⁶

Truth (empirical adequacy) *does* play a role in the examples discussed above: both critics and advocates of the PSC think it captures real taxonomic units in the world, and that it is truthful in this sense. 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 PSC in a particular context of research with concrete aims. For researchers in conservation contexts, such species are too numerous, too small and too vulnerable – they hamper conservation efforts and thus are not significant units for conservation. For the purposes of combating malaria, however, species identified using the PSC *are* significant. The PSC, then, is not selected or criticized on the basis of empirical accuracy or truthfulness, but according to whether the information it captures is significant in particular contexts of scientific practice, where moral and social values related to nature conservation and human health determine significance (cf. Dupré, 1993).

These examples show that non-epistemic values feature in classificatory practices and can play decisive roles there. For the philosophy of classification this means that we need an account of classification and kinds that can accommodate both the epistemic and non-epistemic values scientists employ when choosing the framework for the classification of entities in a domain of study. Available accounts fail to do this. Next, we present an account that does meet this criterion, the GFA, and show (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, 2022; Reydon, 2021) and we can only provide a brief overview here.

⁶ We use Anderson’s notion of ‘significance’ and Kitcher’s notion of ‘significant truths’ without committing to any strict philosophical theory of truth or significance. All Kitcher means by ‘truths’ is well-supported, empirically adequate claims about the world. When talking about ‘significant truths’ his point is that while there are infinitely many possible empirically adequate claims about the world, only a fraction of these are important to us in view of real-world problems. Thus, investigators should not simply aim at finding empirically adequate claims about the world, but rather aim for the important ones. Similarly, Anderson argued that not all questions we could possibly ask are equally important, but scientific research is (and should be) driven by what we think are important questions. We follow Kitcher and Anderson such that ‘truth’ means ‘empirically adequate’ and ‘significant’ means ‘important in light of the interests of a particular group’. In this sense, truth and significance can serve as reasons for theory choice: theories can be more or less empirically adequate, and more or less important to us given our particular interests.

The GFA is motivated primarily by the aim to adequately accommodate the various aims scientists have when classifying a domain of entities into kinds. This is partly in line with traditional views, according to which natural kinds are those kinds standing at the focus of scientific investigations (Bird & Tobin, 2018). However, while available accounts of natural kinds often focus on kinds in the sciences, they tend to have too narrow a view of the role of kinds in the sciences. They tend to focus on specific epistemic aims of natural kinds and the classifications in which they feature, 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: the ones just mentioned, but also achieving a stable reference system for an area of work, achieving a classification that adequately represents evolutionary history, grouping entities to track their observable properties, achieving kinds suitable to control natural phenomena, obtaining kinds for practical purposes such as managing and combating diseases, etc. (Ereshefsky & Reydon, 2015). Often, several aims will be intertwined in one research context, but different contexts (different areas of science, but also 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.

We do not advocate any specific set of legitimate classificatory aims in science, but rather a thorough naturalism according to which a theory of natural kinds and classifications should accommodate *any* classificatory aims that feature successfully in the sciences *and* the success of which is accounted for by the relevant field of science. By focusing on just one or a few epistemic aims, available accounts are insufficiently naturalistic (Ereshefsky & Reydon, 2015). The naturalism we have in mind is the broad claim that when asking which natural kinds should be recognized, philosophers should “rely on what science tells about the world and [...] eschew aprioristic philosophy” (Sklar, 2010: 1121). Accordingly, accounts that single out one or a few aims as defining natural kinds are insufficiently naturalistic as they fail to follow science in the variety of aims posited in successful classifications (Ereshefsky & Reydon, 2015). Most importantly, they ignore non-epistemic aims that play significant roles in scientific classification.

But note that the GFA’s thorough naturalism does not amount to simply taking at face value whatever scientists say about the classifications they use. First of all, we acknowledge that the question which aims and values are admissible in the context of accepting and rejecting scientific claims is a principal issue in the literature on science and values. A general trend in that literature is to consider as acceptable those values that are democratically endorsed, especially by relevant stakeholders (de Melo-Martin & Intemann, 2016). When we say that an account of classification and kinds should be sensitive to the aims of scientists, we assume that those aims are appropriately influenced by democratic deliberation in society and consistent with the values of relevant stakeholders. While this does not provide a strict distinction between acceptable and unacceptable aims and values, it does provide a first-order distinction between projects that involve aims and values that the scientific community widely accepts and those that are widely considered unacceptable (think of projects aimed at promoting eugenics or views of racial

inequality, and the like). This attitude fits the GFA's naturalistic outlook: the GFA is intended as an account of kinds and classifications in science as it *actually is* and not of kinds and classifications in the context of an ideal picture of science.⁷

Furthermore, the fact that this yields only a rough distinction between acceptable and unacceptable aims and values does not constitute a problem for the GFA, as the GFA also incorporates its own normative force in this regard. The GFA does not accept just any classificatory aim that is successfully achieved in some area of science as individuating a natural kind, but only accepts aims for which the relevant field of science explains why the classification that is used succeeds in achieving the aim in focus. Accordingly, the GFA enables a critical engagement with the aims that scientists claim a classification is intended to achieve, and it does not merely defer all judgment to science (especially when it comes to non-epistemic values). As we explain in what follows, the GFA enables such a critical engagement by specifying two conditions that jointly have normative force. The GFA is designed as an account of kinds and classification that captures the various epistemic and non-epistemic reasons scientists have for positing natural kind classifications, but it does not accept just any grouping as a legitimate natural kind. It only accepts natural kinds that are embedded in a well-confirmed theoretical setting, properly informed about non-epistemic values. This makes it 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 providing us with an ideal model of kinds and classification at a distance from scientific practice.

A central element in the GFA's naturalism is the notion of 'classificatory programs' (Ereshefsky & Reydon, 2015; the basic components of this idea were developed in Ereshefsky, 2001). Classificatory programs are those parts of research programs in which classifications are constructed and 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 successful classificatory programs, 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 should be explained as having identified groupings that represent relevant aspects of the world. Which aspects those are is left open in the GFA, and this will differ among research contexts and classificatory programs. As an overarching account of kinds, the GFA says that successfully used

⁷ We take naturalism seriously in another sense too, namely by not a priori restricting the scope of the GFA to those kinds that feature successfully in Western science. The GFA leaves room for kinds in other contexts, such as classifications used in indigenous communities, folk classifications, classifications in non-Western science, etc., to be recognized as natural kinds on a par with kinds in Western science. The GFA thus is intended as an account of classifications as they actually are used successfully in various contexts. The requirements for a grouping to be recognized as a natural kind are the same in all contexts: the kind must serve a purpose, its functionality must be grounded in the world and the classificatory program in which the kind is embedded must provide an account of how the kind's functionality is grounded in the world. While we lack space to develop this issue here, we will say a little more on this below (especially in footnote 10).

scientific kinds represent aspects of the world, but does not prescribe which aspects those must be.⁸

The GFA is a normative account. It puts two constraints on classifications and the kinds featuring 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. A classification is always 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 for us to accept it as a classification that groups entities into natural kinds, it should satisfy those purpose(s). For instance, *if* a classification is posited to facilitate induction, it should satisfy that aim; *if* it is posited to achieve a stable inventory of a particular domain, it should do that; *if* a classification is posited to help promote biological conservation, it should do that.

The *Grounding Condition* requires that the way in which a classification satisfies its purpose(s) should be grounded in the world. The idea is that successful scientific classifications identify groupings that depend on relevant aspects of the world, and not *merely* on our interests and our conceptions of what the world is like. This poses an additional restriction on natural kinds. The Functionality Condition tells us which groupings are *candidates* for being thought of as natural kinds: those groupings that successfully serve the aims they were posited for are candidate natural kinds, because they have proven useful as kinds that can be studied scientifically (which is traditionally seen as a requirement for natural kindhood – see Bird & Tobin, 2018) and used for various purposes. But meeting this criterion is not sufficient for attributing natural kind status to a grouping. It still has to meet the Grounding Condition, which adds that some subset of the set of groupings that meet the Functionality Condition can be thought of as natural kinds *if* the relevant area of science gives us an account of how these kinds are grounded in the world and how this way of grounding enables them to serve their purpose(s).⁹ In sum, the GFA tells us that those groupings that meet both the Functionality and Grounding Conditions should be considered natural kinds.

⁸ In this sense, too, the GFA is less restrictive than other accounts in recognizing natural kinds. For example, on Magnus's account (which is presented as a naturalistic account of natural kinds), natural kinds are those kinds that "scientists are *forced to posit* in order to be scientifically successful in their domain of enquiry" (2012: 47; emphasis added). The GFA, in contrast, does not assume that there is one best way of classifying per domain of inquiry that researchers would be forced to posit for their projects to be successful. On the GFA, kinds are recognized as natural kinds if they perform a role in a research project *and* this role is grounded in the world, thus in principle allowing a large number of kinds to be acknowledged within one and the same research project. In contrast to Magnus' account, the GFA sees researchers as usually having a spectrum of classifications to choose from.

⁹ An anonymous reviewer asked what the GFA requires to be grounded: the kind itself or its functionality? The GFA primarily requires the latter – as we stated above, the Grounding Condition requires that *the way in which a classification satisfies its purpose(s)* should be grounded in the world. Note, though that this seems to imply that the kind itself must also be grounded. It seems strange to say that a kind's functionality can be grounded in the world while the kind itself is not grounded. The opposite, however, *can* hold: a grouping can be grounded while not being particularly functional (e.g., the group of all objects with mass larger than 1 kg is grounded in the world, namely in the fact of the matter that some objects have a mass larger than 1 kg and others do not, but it isn't a functional grouping). This is why the GFA emphasizes that functionality must be grounded. This opens up a deeper metaphysical question, but due to space constraints we cannot address it here but merely point to it.

Note that the two conditions are not independent. The GFA recognizes kinds as natural kinds only if they are jointly met: the kinds in question must successfully achieve the aims for which they were posed *and* the area of science in which they feature must have a well-confirmed theory that explains why these kinds achieve those particular aims. To make this more concrete, let us clarify our use of ‘grounding’ a little bit (but for reasons of space, we cannot address the precise connections between the GFA’s notion of ‘grounding’ and the extensive metaphysical literature on grounding). With ‘grounding’ we invoke the basic point that for a kind to be useful, its functionality must in *some* way be anchored to, or supported by, aspects of the world. As mentioned above, according to the GFA natural kind classifications in part depend on the world and not *only* on our conceptions of what the world is like. The term ‘grounding’ here thus refers to a relation of metaphysical dependence of a kind upon aspects of the world.

By ‘grounding’ we just mean the following: classificatory programs make certain assumptions about the world to construct classifications (i.e., assumptions about connections between the properties of entities that are classified and features of the world that underpin these properties), and a classification can be accepted as classification that groups entities into natural kinds *only* if those assumptions are correct about the world (Ereshefsky & Reydon, 2022). With respect to the question of the correctness of the assumptions underlying a classification, the same philosophical criteria apply as for the question of the correctness of items of scientific knowledge more generally. The GFA does not ask scientists to prove the truth of their assumptions, but rather asks for well-supported assumptions, and for assumptions that can be tested and confirmed in the first place. Here, the onus is on relevant areas of science to explicate what aspects of the world a classification connects to, to explain how these aspects of the world enable the classification to achieve the epistemic and non-epistemic aims that are in focus, and to do so in a way that is testable. In other words, classifications and their grounding must be *defeasible* (Ereshefsky, 2018) – we must be able to test the claims that a classificatory program makes about how kinds are grounded and decide whether these claims are well-supported.

One way of doing this is to require that the classification is embedded in a well-confirmed theory that explains how the aims of a classification (the motivating principles of a classificatory program, which include both epistemic and non-epistemic aims) are connected to the parameters used as the basis of the classification (the sorting principles).¹⁰ In this context, the philosophy of science provides the tools to

¹⁰ But note that this is just *one* way of doing this. While in the present paper our focus is on kinds that feature successfully in science, as noted in footnote 7 the GFA leaves room for kinds outside Western science to be recognized as natural kinds. In such contexts, the GFA requires kinds to be embedded in classificatory programs too, and requires that such programs provide empirically confirmed accounts of how a kind satisfies its purpose(s). In these other contexts such accounts would not be scientific theories in the strict sense of the literature in the philosophy of science (which tends to be focused on Western science), but they *would* have to have some explanatory import with respect to the question how kinds satisfy their purposes. This is not an issue that we can satisfactorily address here, but we want to draw attention to the fact that the GFA does not require all classificatory programs to involve full-blown scientific theories – only scientific classificatory programs have to, while for other contexts parallel requirements hold that should be formulated in different terms. What is generally required is a sufficiently supported account that explicates how the Functionality and Grounding Conditions are met.

distinguish between well-confirmed explanations and insufficiently confirmed ones. Note that this does not introduce a third condition for kinds to be counted as natural kinds. If a kind meets the Functionality Condition and the Grounding Condition, the kind can be counted as a natural kind. But to be able to assess whether those two conditions are met, an account is needed that explicates what enables the kind to perform the function(s) it was posited for. Such an account should be provided by the classificatory program in which the kinds is embedded, and it must be more than just a speculative story. Speculative stories are always possible, such that if the GFA would allow speculative stories, any grouping could count as a natural kind and the GFA would not have normative force. Therefore, a reasonable requirement is that the classificatory program provides a well-supported account of how the Functionality and Grounding Conditions are met (where applying the usual criteria for support for scientific claims is one way of assessing this). If there is no such support, what reasons would we have to accept that the two conditions are indeed met?

Illustrative examples can be found in pseudosciences, as well as in folk taxonomies. For example, the various kinds of phenomena studied in parapsychology, such as clairvoyance, psychokinesis and telepathy, presumably serve the non-epistemic aims of parapsychologists (for example, defining a research community, obtaining funding, etc.), but cannot count as natural kinds due to the lack of a well-confirmed theory (or other sort of account) that plausibly explains how these aims are connected to aspects of the world. Of course, there may be well-confirmed theories in for instance sociology or psychology that explain how parapsychological classifications further the non-epistemic aims of parapsychologists by explicating how the kinds latch onto interests that are present with certain groups in society. But note that the GFA requires that a classificatory program *itself* provides such a theory: the motivating and sorting principles of a classificatory program must be connected by the program in such a way that the motivating principles provide theoretical support for the sorting principles. In the case of parapsychological classification, the requirement thus is that the parapsychological theoretical framework itself explains how the classification promotes the program's aims – i.e., how the classification is grounded in those aspects of the world that parapsychology investigates. For the case of parapsychology, this requirement clearly is not met.

Similarly, while a folk classification of cetaceans as fish may achieve the aims of communities of fishermen (e.g., Dupré, 1993: 29-30; Souza & Begossi, 2007), the GFA does not count such a classification as a classification of organisms into natural kinds as long as the relevant community does not provide a well-confirmed account that connects these aims to aspects of nature. Such an account does not have to be a scientific theory in the strict sense of the literature in the philosophy of science, but it *does* need to be an empirically supported, explanatory account. The classification of cetaceans as fish *does* depend on aspects of the world: animals are grouped together on the basis of their habitat and their similar ways of living in that habitat that make them salient for fishing practices. The category of fish in this context represents a life-form (Souza & Begossi, 2007: 11). But as a category it is simply defined by those properties that are relevant for the aims in view, without any additional theoretical embedding. That is, the classificatory context does not provide any

explanation for how the classification is grounded in aspects of the world – it merely takes every entity that exhibits some non-theoretically related properties as members of the category. Note that this is *not* to say that the classification is wrong – it works well for the specific purposes of its particular context of practice (local fishing practices). But the classification does not identify a natural kind because there is no account (yet) that grounds its functionality in aspects of the world, which is what the GFA requires.

The GFA is not intended to account for all kinds that may be used successfully in any context of the lives of some group of people, but only to account for those kinds that are, or could in principle be, studied scientifically. The paradigms of such kinds are kinds that serve a purpose in scientific contexts (Ereshefsky & Reydon, 2015, 2022), and in this respect, the GFA stands in the contemporary tradition of the philosophy of natural kinds (Bird & Tobin, 2018). This is why the GFA requires that the area of science in which a classification features must have a well-confirmed theory that explains how the epistemic and non-epistemic aims the classification serves are connected to aspects of the world. This also limits the classificatory programs considered in the GFA to those that can be considered potentially acceptable parts of the natural, life, social, and engineering sciences. But, again, note that this does not a priori limit the range of acceptable kinds to those that are studied in contemporary Western science: any folk kind, for example, for which there is an account that explains how it serves its purpose(s) could in principle be studied scientifically, even though at present there is no area of science that actually is concerned with it. In this way the GFA is not a priori limited to being an account of scientific kinds, even if it starts from examining the kinds that feature in our current best science (because that is where the best examples of natural kinds are found). With respect to which classificatory programs are acceptable, the GFA can defer to the tools available to philosophers of science to assess such matters – accounts of theory confirmation, of distinctions between science, and pseudoscience and non-science, of the nature of scientific explanations, and so on.

The GFA does not presuppose any particular way of grounding, but is neutral on this issue. As a naturalistic account of kinds, it allows for different ways in which the functionality of kinds (and kinds themselves) can depend on the world – however, with the proviso that the area of science in which the classification occurs must itself be able to explain how the classification is grounded and how this way of grounding promotes the aims of the classification. This means that satisfying the Functionality and Grounding Conditions is a *local* matter: the conditions should be satisfied in relation to the specific aims of a classificatory project, and those differ among programs. In contrast to other accounts of classification, the GFA does not assume any overarching (set of) aim(s) for all classifications.

This approach fits actual scientific practice. Many species concepts, for instance, were intended for application only to a specific range of organisms (e.g., sexually reproducing organisms, or only birds), rather than throughout all of biodiversity. Accordingly, their success should be measured only with respect to their intended application. Note, too, that the Functionality and Grounding Conditions work together, as the aims of a classificatory program

constrain the possible ways in which a classification can be grounded in the world. Consider again the BSC. Its aim is to posit groups of organisms that are evolutionary units (groups of organisms evolving in tandem) within the research context aiming to explain how speciation occurs and how species remain in existence for extended periods of time (Mallet, 2010, 2013: 682). The BSC satisfies the Functionality Condition *if* it successfully sorts organisms into evolutionary units. It satisfies the Grounding Condition *if* the way it sorts organisms (according to successful interbreeding within a species and reproductive barriers among species) is indeed one way the world causes organisms to cluster in evolutionary units. Both conditions are in fact satisfied, so according to the GFA the BSC identifies natural kinds.¹¹

4 The GFA in scientific practice

4.1 How the GFA handles decisions on species concepts

The two conditions in the GFA can serve as a tool to analyze specific cases. We now deepen our analysis of the examples discussed earlier (Section 2), where we highlighted the role of non-epistemic values in decisions on the acceptance of the PSC. We show how the GFA handles the acceptance of classifications according to non-epistemic values and, by comparisons with other accounts of kinds and classification, how the GFA's machinery works.

First, the conservation case. Three very straightforward criteria in conservation contexts are diagnosability, applicability and practicability. Species must be diagnosable to apply conservation measures to them – that is, we must be able to decide to which species a local population belongs in order to decide whether protective measures should be applied to it. This is why cryptic species have long been a problem for species conservation – a problem that is increasingly being solved by improved methods for fine-grained genetic analysis (Bickford et al., 2007; Chenuil et al., 2019). Applicability is the requirement that a species concept can be applied throughout as many domains of biodiversity as possible and does not only apply to, for example, animals. In this respect, the BSC does not meet this criterion well, as it only applies to sexually reproducing organisms and thus only to a minority of biodiversity. Wide applicability of a species concept is important to enable comparisons between different species with respect to their conservation status: if a species of plant and a species of animal are individuated using different species concepts, it is difficult to decide which should be prioritized in conservation contexts because the comparison is between two different kinds of things. Practicability is a criterion that follows from

¹¹ One might wonder if this conflicts with the metaphysical claim that species are individuals (Hull, 1978). However, the GFA is neutral on such “deep” metaphysical issues and simply says natural kinds are those groupings that feature successfully in science, where success is explained using the Functionality and Grounding Conditions. The GFA does not aim at determining which metaphysical categories the units of a classification belong to – metaphysically they can be individuals, classes, both, or neither (cf. Boyd, 1999: 163; Brigandt, 2009: 78) and still be natural kinds in the sense of the GFA.

the fact that the actual implementation of conservation measures costs resources. If a species concept is used that recognizes too many species that are threatened with extinction, for example, there may well be insufficient funds and human power available to protect them all. Thus, a situation may occur in which a species concept that allows a fine-grained diagnosing of species yields many more species than we have capabilities to conserve – i.e., a conflict between diagnosability and practicability.

Participants in the debate generally agree that the PSC is favored over the BSC and other species concepts if *only* diagnosability and wide applicability were considered. The PSC generally yields very fine-grained diagnosable groupings, because it uses unique traits (character states that originated in one branch of the Tree of Life) in combination with common descent to group organisms (Reydon & Kunz, 2019).¹² The uniqueness of these traits guarantees diagnosability of the ensuing groups. The PSC often meets the diagnosability requirement better than the BSC, as groups defined by reproductive isolation are very difficult to diagnose in the wild or the laboratory (Agapow et al., 2004: 163). Also, it is applicable throughout the whole of biodiversity, setting it apart from other species concepts (*ibid.*), whereas the BSC is applicable only to sexually reproducing organisms. Thus, the PSC performs better than competitor concepts on two important epistemic values – diagnosability and applicability. This notwithstanding, conservation biologists criticize the PSC for non-epistemic reasons. As shown in Section 2, using the PSC tends to yield too many species, and it yields species that are too small and thus become threatened too easily, thus increasing the burden on limited conservation resources. Crucially, in the debate we find a trade-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 to override the former.

Philosophical accounts of kinds and classification thus should not just make room for non-epistemic values to play some role in classifications (e.g., as factors operating in the background), but should explicitly allow the possibility that non-epistemic values override epistemic values. Therefore, an adequate account of kinds and classification cannot a priori prioritize epistemic values over non-epistemic values (or, for that matter, non-epistemic values over epistemic values), but must be open to either way of prioritizing and follow how researchers *actually* prioritize values in practice. While often epistemic values will be most prominent, in many other cases non-epistemic values override epistemic values (and we discuss examples in which this is the case). Given their strong (and sometimes even exclusive) focus on epistemic values, available accounts of natural kinds fail to meet this condition. In Section 4.2, we show how the GFA performs better in this respect in comparison to some other accounts.

Before getting to that issue, we want to address another aspect of the GFA, namely that functionality and grounding can come apart. Consider as an example the question whether gene-edited organisms (GEOs) should ontologically be counted

¹² The PSC exists in several versions. The version quoted in Section 2 is sometimes called the “diagnosable version” (Maiden, 1997: 405), but the point we make holds for all versions.

as genetically modified organisms (GMOs) or as a separate category.¹³ This is an important question in relation to the legal regulation of agriculture and food production (Wasmer, 2019; Lohse et al., 2020). While many consumer groups and environmental organizations argue that GEOs should be seen as GMOs, many plant scientists argue that GEOs are not GMOs, as GEOs do not include transgenes and could in principle have come into being due to natural spontaneous mutations. Behind the latter arguments often are non-epistemic aims: excluding GEOs from the category of GMOs makes research less complicated, among other things, because no special permissions are required and industry will have a greater incentive to fund research on GEOs when there will not be any obstacles to marketing GEO-derived foodstuffs.

While a classification that separates GEOs and GMOs as distinct kinds thus meets the non-epistemic aims of the involved scientists, this is insufficient to count the groups ‘GEO’ and ‘GMO’ as natural kinds under the GFA. This is because the achievement of the specific aims in question (making research easier, having less restrictions on research, generating incentives for industry funding) is not due to any specific way in which the two groups would be grounded in the world. While GEOs differ from GMOs in the lack of transgenes, whether or not this difference is relevant for making research easier fully depends on how we legally treat this difference. For one, transgenic organisms can also come into being in natural ways (by way of horizontal gene transfer), such that the aspect of natural vs. artificial origins does not ground a distinction between GEOs and GMOs as kinds. Furthermore, GEOs *do* have artificial origins, even though similar organisms *could have* also occurred naturally. Research on GEOs thus is not made easier because of any fact of the matter that GEOs are, or are not, GMOs – the entire effect of the classification is due to *how we legally treat* GEOs and GMOs, and not to any difference in nature between GEOs and GMOs that is relevant from the perspective of scientific theory. Both GEOs and GMOs result from technological interventions in organisms that change their genetic makeup – the difference is that for GMOs these interventions are perceived as involving genetic material that is “foreign” to the species, leading to differences in legal treatment. In both cases, the sequence of As, Cs, Ts and Gs of the organisms involved is changed – whether or not GEOs should be seen as GMOs (and thus whether or not research would be made less complicated) is purely a matter of how we legally treat these changes. That is, while the Functionality Condition is met in this case, from the perspective of those plant scientists that want to distinguish GEOs and GMOs, the Grounding Condition is not.¹⁴ The groupings ‘GEO’ and ‘GMO’ satisfy the Functionality Condition, but the distinction in their functions is not grounded in the world

¹³ This case was posed to us as a possible counterexample by an anonymous reviewer.

¹⁴ Note that the situation is different for the activists’ grouping of GEOs as GMOs. If their aim, for example, is to mitigate risks for human health and for the environment that are due to genetic technologies (such as errors in gene editing or genetic modification processes), grouping GEOs and GMOs together promotes this aim. Also, the grouping of GEOs as GMOs is grounded (by representing a category of processes), and the relevant areas of science and technology explain why the grouping promotes the aim (namely by highlighting a group of entities that is made in technological processes that carry a certain kind of risk). With respect to the specific purpose of promoting human health, functionality and grounding do *not* come apart and accordingly the GFA tells us that the grouping together of GMOs and GEOs *in the context of this particular purpose* yields a natural kind.

(but only legally defined) and hence they should not be thought of as distinct natural kinds. This example shows how the GFA does not simply defer to science but enables a critical and normative attitude towards scientific classifications: its normativity derives from the Functionality and Grounding Conditions working in tandem. When functionality and grounding come apart, the GFA tells us that such functional groups – notwithstanding their functionality – should not be counted as natural kinds.

This shows how the GFA's thorough naturalism explicitly does *not* involve a view that any aim is as good as any other, such that scientists should construct classifications to further any sort of aim (e.g., political aims, aims related to obtaining funding or improving prestige, etc.) and as philosophers we would have to accept all these classifications as involving natural kinds. The GFA does defer to science with respect to the aims for which classifications are constructed, with the caveat that such aims should be informed by values that are democratically endorsed, especially by relevant stakeholders (see Section 3). But at the same time the GFA poses the requirement that a classification actually succeeds in achieving these aims *and* that the relevant field of science can explain why this is the case.¹⁵ The former requirement is embodied in the Functionality Condition, the latter in the Grounding Condition. In the first step of the analysis, the Functionality Condition is used to identify successful kinds in science that are *candidates* for being attributed natural kind status: success in achieving the aims for which a kind was posed thus is a necessary but not sufficient condition for accepting it as a natural kind. In the second step this group of candidates is narrowed down, using the Grounding Condition 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 works to identify natural kinds *and* explain their successful use in scientific research.

This can also be seen in our example of the classification of organisms into species. The GFA reconstructs the debate on species concepts in conservation contexts as follows. As it meets the applicable epistemic requirements, the PSC seems to meet the GFA's Functionality Condition, such that the groups based on it seem to be putative natural kinds (and we would have to invoke the Grounding Condition to decide whether they actually should be given natural kind status). Adding non-epistemic values into the picture, we see why the Functionality Condition is in fact not met by the PSC. While diagnosability furthers the aims of conservation to *some* extent (as a minimal level of diagnosability is required for conservation efforts to be effective), it yields too many diagnosable groups. In terms of the GFA, the opponents of the PSC (Frankham et al., 2012; Zachos et al., 2013) argue that it fails to meet the Functionality Condition, as it does not yield groupings that can feature in successful conservation efforts.

Had we only considered epistemic aims in our analysis, the conclusion would have been that the PSC meets the Functionality Condition and criticism of the PSC was misguided. Considering both epistemic and non-epistemic aims allows us to give an appropriate reconstruction of the debate and fundamentally changes the conclusion regarding criticism of the PSC. Note that in this analysis, the Functionality

¹⁵ Again, this latter requirement is not met in the case of GEOs and GMOs. The relevant field of science does not explain why GEOs and GMOs are treated differently from a legal perspective.

Condition exerts some 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 subordinate epistemic aims are diagnosability and broad applicability. This allows us to see that authors rightly criticize the PSC (if their empirical claims are correct), because they emphasize the actual aims of the program for which the classification is intended.¹⁶ As the Functionality Condition is not met by the PSC in this case, the Grounding Condition doesn't come into play, as the GFA tells us to first test whether the Functionality Condition is met and second whether the Grounding Condition is met.

We recognize, however, that talking about “the actual aims” of a research context is a vexing issue.¹⁷ This is not a specific problem for our account, but a more general problem in the literature on science and values. We cannot solve it here, but we want to highlight how we deal with it. First, by examining cases from uncontroversial areas of science (rather than projects that promote controversial aims), we can show how our account works in several actual cases of scientific practice. We do this in the present paper as well as in other work (Ereshefsky & Reydon, 2022; Reydon, 2021). Second, we acknowledge that it is possible for disagreement to exist within one research context on the aims of research and relevant values, such that it is impossible to identify *the* aims of research for this particular context. But there is a way to treat such cases, we suggest, namely by thinking of a context in which disagreement on aims exists as encompassing multiple distinct classificatory programs, each pursuing its own aims and coexisting with the other classificatory programs within the larger research context. Our notion of classificatory programs makes this possible, as a classificatory program is *defined* in part by its aims (i.e., its motivating principles), such that disagreement on aims exists *between* programs. For example, the various competing approaches in biological systematics (see Hull, 1988) can be considered different classificatory programs operating within the same larger research context but with different aims in view and, as Hull described in detail, heated disagreements between programs. Classificatory programs as we understand them thus are very localized and because the GFA assesses the kinds that are posed in one classificatory program at a time and does not perform comparative assessments between different classificatory programs, it is able to avoid disagreements about aims.

Now consider 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

¹⁶ One might ask if this case is not just a common case of classificatory pluralism, i.e., a case in which different epistemic values support different choices of species concepts and non-epistemic values merely provide additional motivation when choosing among concepts. However, the case of conservation we are discussing is *not* such a case. A crucial aspect of this case is that epistemic values and non-epistemic values *conflict* as to which is the best concept for the job, and researchers chose to let non-epistemic values override epistemic values. It would thus not be possible for researchers to make the choice they made without strongly relying on non-epistemic values as the main factor. This sort of case is unaccounted for in the literature on classificatory pluralism. The GFA, however, does account for cases like this – which is why it is needed here.

¹⁷ This clarification is in response to requests from two anonymous reviewers, to whom we are indebted for raising this issue.

different conclusion here. Note first that (as discussed above) Attenborough (2015) argues that when it comes to the prevention and eradication of malaria the capability of diagnosing cryptic species favors the PSC. Using the PSC, we can achieve a more fine-grained taxonomy of species in the genus *Anopheles*, enabling more differentiated associations between the local presence of a particular type of mosquito and connected risks for human health than on the basis of traditionally recognized species. As Attenborough shows (2015: 144; Attenborough's Table 7.2), using the PSC allows us to split two recognized species of *Anopheles* into seven species differing in their ranges of occurrence, population densities, and the extents to which their organisms can carry malaria parasites and to which they prey on humans or rather on non-human animals.¹⁸ Here, too, diagnosability is what counts. But because of the aforementioned differences between the species recognized using the PSC, the epistemic aim of diagnosability here aligns with the non-epistemic aim of the promotion of human health in regions where malaria occurs.

Here, too, the Functionality Condition exerts normative force, telling us to prioritize the actual aims of the research program under consideration (malaria epidemiology, aiming to promote human health). Following this guideline, we see that the PSC fits the actual aims of this specific context better than other species concepts (assuming the scientific claims in Attenborough's paper are correct). The Grounding Condition adds to this by requiring that the groups identified by the PSC represent aspects of the world relevant to the aims in malaria research, i.e., shared traits due to common descent that pertain to whether or not mosquitoes can carry malaria parasites and whether or not they prey on humans. The PSC rests on common descent as well as unique traits, and as such highlights those traits that explain why different groups of mosquitoes have different relevance when it comes to malaria research and prevention. The Grounding Condition thus is met and the available body of biomedical knowledge explains how it is met. The analysis 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 for which the classification is intended. Here, too, non-epistemic values are decisive and play a crucial role in the choice made by the researchers.

A legitimate worry (voiced by one of the anonymous reviewers of this paper) would be whether the assumption is correct that there is general agreement about the aims of research among the researchers in a particular context of research. In the case discussed above, for example, a question would be whether the local aims of research on malaria in Pacific regions and of research on malaria in Africa are sufficiently aligned for the same classification to serve both sets of aims. While we agree that this may be an issue, it is not an issue for the GFA to deal with. As explained above, the GFA works by taking the aims as explicated by researchers working in a particular classificatory program and assessing whether the classification that is proposed in that program successfully meets these aims and the program is able to explain how it does so. When researchers disagree on research aims, the GFA treats this situation as a disagreement between classificatory programs that either compete

¹⁸ *A. punctulatus* (in Pacific regions) splits into *A. punctulatus*, *A. farauti*, *A. hinesorum*, and *A. irenicus*. *A. gambiae* (in Africa) splits into *A. gambiae*, *A. arabiensis*, and *A. quadriannulatus*.

within the same context of research or apply to different subareas of research in the same context. In the malaria case, for example, researchers might find that the promotion of human health with respect to malaria in Pacific regions requires a different classification than in Africa. As a naturalistic account of kinds, the GFA would then follow the researcher's findings and assess both classifications separately, asking whether each meets the specific aims it was set for. In the case of competing programs, for example one group of researchers advocating one classification for all malaria research and another group advocating different classifications for different geographic regions, the GFA would be able to exert some normative force. The GFA could be used to assess the competing programs for whether they group organisms into natural kinds. If one of them does not, this would indicate a problem for that program with respect to the grounding of its classification in the world. If all do, this would indicate a need for the different programs to examine how their classifications could be related to each other, for example in a hierarchical manner.

4.2 The GFA in comparison

In both examples involving species we see how the GFA can make sense of the scientists' reasoning. 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 local classificatory programs as the basis for analysis and being fundamentally open to the possibility that non-epistemic aims override epistemic aims. In this respect the GFA contrasts strongly with other prominent accounts. To clarify the contrast, we briefly look at Khalidi's (2013, 2018) "causal nodes" account, Slater's (2013, 2015) Stable Property Cluster account, and Boyd's (1991, 1999) HPC account as examples.

On Khalidi's account, legitimate natural kinds are groupings of entities that represent nodes in causal networks (i.e., in the causal structure of the world). Khalidi (2013, Section 4.7) holds that classifications should represent such nodes and excludes non-epistemic values from playing a role in determining natural kind status. On Slater's account, legitimate natural kinds are groupings of entities that represent stable patterns that we find in the world – i.e., stably recurring patterns of similarity between entities. Because patterns can be stable to higher and lesser degrees, on Slater's account groups can be attributed kind status to higher and lesser degrees. Slater expresses this with his notion of "natural kindness" – being a natural kind (natural kindness) is not a yes-or-no matter but comes in degrees. Khalidi and Slater use *global* criteria that all kinds in all contexts must meet, and that do not differentiate between kinds that meet them. On Khalidi's account any classification that represents 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. Similarly, 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 and the degree of stability satisfies the norms of a discipline, it should be given natural kind status. What the two accounts lack are

filters that would allow us to distinguish more from less *significant* causal nodes and more from less *significant* stable patterns, respectively.

On Slater's account, for instance, non-epistemic interests do not play a role in determining whether a stable pattern *is* a natural kind or not – rather, epistemic interests determine which of the many available stable patterns (or: kinds) researchers focus on and which they disregard. A fundamental difference between Slater's account and the GFA is that non-epistemic values and interests are *internal* to the GFA but *external* to Slater's account, as for him they do not play a role in determining whether or not a group should be attributed natural kind status. Natural kind status is attributed solely on the basis of being a stable cluster of properties for the epistemic purposes of a discipline (Slater, 2015: 396). Non-epistemic values only play a role in selecting which kinds to focus on. On the GFA, in contrast, interests are *co-constitutive* of what it is to be a natural kind, as kinds must meet the Functionality Condition. In this sense, non-epistemic values are internal to the GFA. The GFA thus can reject patterns that fail non-epistemic interests and, conversely, allow groups of interest to researchers without being connected to any stable property cluster. By internalizing non-epistemic values, the GFA (in contrast to Slater's account) can highlight factors that make the difference for scientists in the examples we discussed.

Note 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 arguing against the use of the PSC in conservation contexts. After all, the PSC works perfectly well when it comes to identifying causal nodes (here inheritance from a common ancestor can be interpreted as a causal node – see Ludwig, 2018: 44-45) or stable patterns (shared traits that uniquely define a group, even if not all members actually exhibit them). Also, we would have to conclude that while Attenborough is right to prefer the PSC 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 PSC is more powerful than competing concepts when it comes to identifying stable patterns (as it recognizes more finely grained patterns). On Khalidi's and Slater's accounts of classification, Attenborough should not have preferred the PSC on the basis of whether it is able to individuate groups of mosquitos that are important for promoting human health.

The crucial problem with both Khalidi's and Slater's accounts, then, is that they miss the aims that in actual research contexts are set by the non-epistemic values endorsed by the community of researchers. As we have seen, these non-epistemic aims do not necessarily 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.¹⁹ But in the examples we discussed,

¹⁹ 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.

the non-epistemic aims matter the most and are taken by researchers to override epistemic values. One important thing these examples of actual scientific practice show 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 featuring in them is acknowledged in some of the available accounts. Boyd's account, for example, explicitly conceives of kinds as relative to disciplinary matrices (Boyd, 1999: 148): 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 classifications according to the aims of local research contexts and programs, which may be a specific school of thought or a concrete practical project within a small subdiscipline.²⁰

Furthermore, while on Boyd's account kinds depend on nature as well as on human classificatory activities (his "bicameralism thesis"), there are important dissimilarities to the GFA. On Boyd's account, only homeostatically supported property clusters are recognized as the aspects of nature on which kinds depend. The GFA is much less restrictive in this respect, as grounding can be realized in many ways and the GFA does not presuppose that there is only one way of depending on the world that determines kindhood. Also, 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 (including non-epistemic aims). So, on both counts the GFA is more naturalistic than Boyd's account.²¹

While we have discussed several accounts that strongly contrast with the GFA, we acknowledge a general trend in the literature on natural kinds with which the GFA aligns. As mentioned at the beginning of this paper, the role of non-epistemic values

²⁰ See Brigandt (2009) for a perspective on kinds that is also more fine-grained and local than Boyd's. Brigandt's view differs from the GFA in important respects, though, most prominently in its strong focus on the epistemological roles of classificatory concepts without an equally strong consideration on how kinds are grounded in the world (which the GFA explicitly includes).

²¹ For more extensive criticism of Boyd's account, see Ereshefsky & Reydon (2015).

in classificatory context has been widely acknowledged and one could say that there is a growing movement among authors on this topic to allow non-epistemic values to contribute to the setting of classificatory aims (e.g., Dupré, 1993; Magnus, 2012; Ludwig, 2014, 2016; Brigandt, 2015; Conix, 2018). Our aim here was to contribute to that movement by proposing a general approach to assessing under which conditions epistemic as well as non-epistemic values can function to individuate natural kinds in scientific contexts. Furthermore, 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, enabling us to reconstruct the decisions that scientists make better than other accounts. These aspects are: that non-epistemic values play important roles in classifications, and that non-epistemic values may 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 philosophy of science. We have tried to show why this is a lacuna in philosophical research – that non-epistemic values sometimes play decisive roles in classificatory practice – and presented an account of kinds and classification that fills this lacuna.

Our account is thoroughly naturalistic and practice-oriented. It does not intrinsically prioritize epistemic values over non-epistemic values, nor does it assume the opposite prioritization, because in some cases researchers will take epistemic values to override non-epistemic values, whereas in other cases this will be the opposite. As available accounts of natural kinds 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 consistent with actual scientific practice. Our account, the GFA, aims to change the focus of philosophical thinking about scientific kinds and classification in this respect.

It is a general account of natural kinds in that it covers all commonly discussed cases in which researchers pose kinds to achieve epistemic aims and also allows philosophers to properly handle cases in which non-epistemic values and aims play a prominent role or override epistemic aims and values. Such cases are plentiful in the life sciences, the biomedical sciences and the social sciences (e.g., Dupré, 1993; Hacking, 1995, 2007) and to some extent also occur in the physical sciences (Anderson, 1995), and have led to a profound rethinking by philosophers of science of the nature of natural kinds. While we have only discussed a few case studies from the life sciences, there is nothing in the GFA that would intrinsically limit its application to that area of science or to classifications outside the scope of Western science. The GFA's core elements – the notion of classificatory programs, the Functionality Condition and the Grounding Condition – apply throughout all areas of research. And they presumably apply outside research contexts too, as motivating and sorting principles feature in everyday classificatory contexts too. Also, the GFA does not encompass any assumptions regarding the question in which areas of research

non-epistemic values and aims play a role with respect to classifications, but can straightforwardly be applied in any area in which non-epistemic values and aims turn out to play a role. This is another aspect of the naturalism embodied in the GFA: when applying it to assess a classificatory program in some area of research, one can simply take the role of non-epistemic aims and values (or the lack of such a role) at face value and apply the GFA whatever this role might be.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest None applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

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.
- Bickford, D., Lohman, D. J., Sodhi, N. S., Ng, P. K. L., Meier, R., Winker, K., Ingram, K. K., & Das, I. (2007). Cryptic species as a window on diversity and conservation. *Trends in Ecology and Evolution*, 22, 148–155.
- 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.
- Boyd, R. N. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 141–185). MIT Press.
- Chenuil, A., Cahill, A.E., Délémonty, N., Du Saillant du Luc, E. & Fanton, H. (2019). 'Problems and questions posed by cryptic species: A framework to guide future studies', in: Casetta, E., Marques da Silva, J. & Vecchi, D. (Eds): *From Assessing to Conserving Biodiversity: Conceptual and Practical Challenges*, Cham: Springer, pp. 77–106.
- Brigandt, I. (2009). Natural kinds in evolution and systematics: Metaphysical and epistemological considerations. *Acta Biotheoretica*, 57, 77–97.
- Brigandt, I. (2015). Social values influence the adequacy conditions of scientific theories: Beyond inductive risk. *Canadian Journal of Philosophy*, 45, 326–356.
- Conix, S. (2017): *Radical Pluralism, Ontological Underdetermination, and the Role of Values in Species Classification*, unpublished PhD dissertation, University of Cambridge.

- Conix, S. (2018). Values, regulation, and species delimitation. *Zootaxa*, 4415, 390–392.
- 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.
- Cracraft, J. (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 Melo-Martin, I., & Intemann, K. (2016). The risk of using inductive risk to challenge the value-free ideal. *Philosophy of Science*, 83, 500–520.
- 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.
- Douglas, H. E. (2009). *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- Douglas, H. E. (2016). Values in science. In P. Humphreys (Ed.), *The Oxford Handbook of Philosophy of Science* (pp. 609–630). Oxford University Press.
- Douglas, H. E. (2017). Why inductive risk requires values in science. In K. C. Elliott & D. Steel (Eds.), *Current Controversies in Values and Science* (pp. 81–93). Routledge.
- Dupré, J. A. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Harvard University Press.
- Elliott, K. C. (2017). *A Tapestry of Values: An Introduction to Values in Science*. 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*. Routledge.
- Ereshefsky, M. (2018). Natural kinds, mind independence, and defeasibility. *Philosophy of Science*, 85, 845–856.
- Ereshefsky, M., & Reydon, T. A. C. (2015). Scientific kinds. *Philosophical Studies*, 172, 969–986.
- Ereshefsky, M., & Reydon, T. A. C. (2022). The Grounded Functionality Account of natural kinds. In W. Bausman, J. Baxter, & O. Lean (Eds.), *From Biological Practice to Scientific Metaphysics*. 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.
- Griffiths, P. E. (1997). *What Emotions Really Are: The Problem of Psychological Categories*. University of Chicago Press.
- Griffiths, P. E. (2004). Emotions as natural and normative kinds. *Philosophy of Science*, 71, 901–911.
- Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 351–383). Clarendon Press.
- Hacking, I. (2007). Kinds of people: Moving targets. *Proceedings of the British Academy*, 151, 285–318.
- Haslam, N. (2002a). Kinds of kinds: A conceptual taxonomy of psychiatric categories. *Philosophy, Psychiatry & Psychology*, 9, 203–217.
- Haslam, N. (2002b). Practical, functional, and natural kinds. *Philosophy, Psychiatry & Psychology*, 9, 237–241.
- Hull, D. L. (1978). A matter of individuality. *Philosophy of Science*, 45, 335–360.
- Hull, D. L. (1988). *Science As a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. University of Chicago Press.
- Intemann, K. (2001). Science and values: Are value judgments always irrelevant to the justification of scientific claims? *Philosophy of Science*, 68, S506–S518.
- Intemann, K. (2005). Feminism, underdetermination, and values in science. *Philosophy of Science*, 72, 1001–1012.
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science*, 5, 217–232.
- 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 University Press.
- Khalidi, M. A. (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*. Oxford University Press.
- Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford University Press.
- Lohse, S., Wasmer, M. S., & Reydon, T. A. C. (2020). Integrating philosophy of science into research on ethical, legal and social issues in the life sciences. *Perspectives on Science*, 28, 700–736.
- Longino, H. E. (1983). Beyond “bad science”: Skeptical reflections on the value-freedom of scientific inquiry. *Science Technology & Human Values*, 8, 7–17.
- Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- Longino, H. E. (1996). Cognitive and non-cognitive values in science: Rethinking the Dichotomy. In L. Hankinson Nelson & J. Nelson (Eds.), *Feminism, Science, and the Philosophy of Science* (pp. 39–58). Kluwer.
- Ludwig, D. (2014). Disagreement in scientific ontologies. *Journal for General Philosophy of Science*, 45, 119–131.
- Ludwig, D. (2016). Ontological choices and the value-free ideal. *Erkenntnis*, 81, 1253–1272.
- Ludwig, D. (2018). Letting go of natural kind: Toward a multidimensional framework of nonarbitrary classification. *Philosophy of Science*, 85, 31–52.
- Magnus, P. D. (2012). *Scientific Enquiry and Natural Kinds: From Planets to Mallards*. Palgrave Macmillan.
- Mallet, J. (2010). Group selection and the development of the BSC. *Philosophical Transactions of the Royal Society B*, 365, 1853–1863.
- Mallet, J. (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 M. F. Claridge, H. A. Dawah, & M. R. Wilson (Eds.), *Species: The Units of Biodiversity* (pp. 381–424). Chapman & Hall.
- 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 C. E. Kendig (Ed.), *Natural Kinds and Classification in Scientific Practice* (pp. 59–73). Routledge.
- Reydon, T. A. C. (2021). Functional kinds and the metaphysics of functional levels: In what sense are functionally defined kinds and levels nonarbitrary? In D. S. Brooks, J. DiFrisco, & W. C. Wimsett (Eds.), *Levels of Organization in the Biological Sciences* (pp. 261–277). MIT 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 D. L. Hull, M. Forbes, & K. Okruhlik (Eds.), *PSA 1992: Proceedings of the Biennial Meeting of the Philosophy of Science Association (Volume One: Contributed Papers)* (pp. 13–22). Philosophy of Science Association.
- Rooney, P. (2017). The borderlands between epistemic and non-epistemic values. In K. C. Elliott & D. Steel (Eds.), *Current Controversies in Values and Science* (pp. 31–46). Routledge.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20, 1–6.
- Sklar, L. (2010). I’d love to be a naturalist – If only I knew what naturalism was. *Philosophy of Science*, 77, 1121–1137.
- Slater, M.H. (2013). *Are Species Real? An Essay on the Metaphysics of Species*, Houndmills, Basingstoke: Palgrave Macmillan.
- Slater, M. H. (2015). Natural kindness. *British Journal for the Philosophy of Science*, 66, 375–411.
- Slater, M. H. (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.
- Souza, S. P., & Begossi, A. (2007). Whales, dolphins or fishes? The ethnotaxonomy of cetaceans in São Sebastião, Brazil. *Journal of Ethnobiology and Ethnomedicine*, 3, 9.

- Wasmer, M. (2019). Roads forward for European GMO policy – Uncertainties in wake of ECJ judgment have to be mitigated by regulatory reform. *Frontiers in Bioengineering and Biotechnology*, 7, 132.
- Zachar, P. (2014). *A Metaphysics of Psychopathology*. MIT Press.
- Zachar, P. (2015). Psychiatric disorders: Natural kinds made by the world or practical kinds made by us? *World Psychiatry*, 14, 288–290.
- Zachar, P., & Potter, N. N. (2010). Personality disorders: Moral or medical kinds—Or both? *Philosophy, Psychiatry, & Psychology*, 17, 101–117.
- Zachos, F. E., Apollonio, M., Bärman, E. V., et al. (2013). Species inflation and taxonomic artefacts—A critical comment on recent trends in mammalian classification. *Mammalian Biology*, 78, 1–6.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.